arXiv:1702.06488v3 [stat.CO] 12 Apr 2017

Distributed Estimation of Principal Eigenspaces*

Jianqing Fan, Dong Wang, Kaizheng Wang, Ziwei Zhu

Department of Operations Research and Financial Engineering

Princeton University

Abstract

Principal component analysis (PCA) is fundamental to statistical machine learning. It extracts latent principal factors that contribute to the most variation of the data. When data are stored across multiple machines, however, communication cost can prohibit the computation of PCA in a central location and distributed algorithms for PCA are thus needed. This paper proposes and studies a distributed PCA algorithm: each node machine computes the top K eigenvectors and transmits them to the central server; the central server then aggregates the information from all the node machines and conducts a PCA based on the aggregated information. We investigate the bias and variance for the resulting distributed estimator of the top K eigenvectors. In particular, we show that for distributions with symmetric innovation, the distributed PCA is "unbiased". We derive the rate of convergence for distributed PCA estimators, which depends explicitly on the effective rank of covariance, eigen-gap, and the number of machines. We show that when the number of machines is not unreasonably large, the distributed PCA performs as well as the whole sample PCA, even without full access of whole data. The theoretical results are verified by an extensive simulation study. We also extend our analysis to the heterogeneous case where the population covariance matrices are different across local machines but share similar top eigen-structures.

Keywords: Distributed Learning, PCA, One-Shot Approach, Communication Efficient, Heterogeneity.

*The research was supported by NSF grants DMS-1206464 and DMS-1406266 and NIH grant R01-GM072611-12.

1 Introduction

Principal component analysis (PCA) (Pearson, 1901; Hotelling, 1933) is one of the most fundamental tools in statistical machine learning. The past century has witnessed great efforts on establishing consistency and asymptotic distribution of empirical eigenvalues and eigenvectors. The early classical work of Anderson (1963) studied the asymptotic normality of eigenvalues and eigenvectors of sample covariances from multivariate Gaussian distribution with dimension d fixed and sample size n going to infinity. Recent focus moves on to the high-dimensional regimes, i.e., both n and d go to infinity. A partial list of such literatures are Johnstone (2001); Baik et al. (2005); Paul (2007); Johnstone and Lu (2009); Jung and Marron (2009); Onatski (2012); Shen et al. (2016); Wang and Fan (2017). As demonstrated by these papers, asymptotic behaviors of empirical eigenvalues and eigenvectors depend on the scaling of n, d and also the spikiness of the covariance. When $n \ll d$, the empirical leading eigenvector $\widehat{\mathbf{v}}_1$ is inconsistent in estimating the true top eigenvector $\mathbf{v}_1$ unless the top eigenvalue λ_1 diverges fast. This phenomenon inspires another line of research on sparse PCA where certain sparsity on top eigenvectors is imposed to overcome the curse of dimensionality; see e.g., Johnstone and Lu (2009); Vu and Lei (2012); Shen et al. (2013); Cai et al. (2013). Besides the asymptotic study, there are also correspondent non-asymptotic results on PCA, for example, Nadler (2008); Johnstone and Lu (2009); Reiss and Wahl (2016).

With rapid developments of information and technology, massive datasets are now ubiquitous. Statistical analysis such as regression analysis or PCA on such enormous data is unprecedentedly desirable. However, large datasets are usually scattered across distant places such that to fuse or aggregate them is extremely difficult due to communication cost. Consider giant IT companies that collect data simultaneously from places all around the world. Constraints on communication budget and network bandwidth make it nearly impossible to aggregate and maintain global data in a single data center. Another example is that health records are scattered across many hospitals. It is hard to process the data in a central location due to privacy and ownership concerns. To resolve these issues, efforts have been made to exploiting distributed computing architectures and developing distributed estimators or testing statistics based on data scattered around different locations. A typical distributed statistical method first calculates local statistics based on each sub-dataset and then combines all the subsample-based statistics to produce an aggregated statistic. Such distributed

methods fully adapt to the parallel data collection procedures and thus significantly reduce the communication cost. Many distributed regression methods follow this fashion (Zhang et al., 2013; Chen and Xie, 2014; Battey et al., 2015; Lee et al., 2015).

Among all the efforts towards creating accurate and efficient distributed statistical methods, there has been rapid advancement on distributed PCA over the past two decades. Unlike the traditional PCA where we have the complete data matrix $\mathbf{X} \in \mathbb{R}^{N\times d}$ with d features of N samples at one place, the distributed PCA needs to handle data that are partitioned and stored across multiple servers. There are two data partition regimes: "horizontal" and "vertical". In the horizontal partition regime, each server contains all the features of a subset of subjects, while in the vertical partition regime, each server has a subset of features of all the subjects. To conduct distributed PCA in the horizontal regime, Qu et al. (2002) proposes that each server computes several top eigenvalues and eigenvectors on its local data and then sends them to the central server that aggregates the information together. Yet there is no theoretical guarantee on the approximation error of the proposed algorithm. Liang et al. (2014) and Boutsidis et al. (2015) aim to find a good rank-K approximation $\widehat{\mathbf{X}}$ of $\mathbf{X}$. To assess the approximation quality, they compare $\|\widehat{\mathbf{X}}-\mathbf{X}\|_F$ against $\min_{rank(\mathbf{B})\leq K}\|\mathbf{B}-\mathbf{X}\|_F$ and study the excess risk. For the distributed PCA in the vertical data partition regime, there is also a great amount of literature, for example, Kargupta et al. (2001), Li et al. (2011), Bertrand and Moonen (2014), Schizas and Aduroja (2015), etc. This line of research is often motivated from sensor networks and signal processing where the vertically partitioned data are common. Our work focuses on the horizontal partition regime, i.e., we have partitions over the samples rather than the features.

Despite these achievements, very few papers establish rigorous statistical error analysis of the proposed distributed PCA methods. To our best knowledge, the only works that provide statistical analysis so far are El Karoui and d'Aspremont (2010) and Chen et al. (2016). To estimate the leading singular vectors of a large target matrix, both of them propose to aggregate singular vectors of multiple random approximations of the original matrix. El Karoui and d'Aspremont (2010) adopts sparse approximation of the matrix by sampling the entries, while Chen et al. (2016) use Gaussian random sketches. The works are related to ours, since we can perceive sub-datasets in the distributed PCA problem as random approximations. However, our analysis is more general, since it does not rely on any matrix incoherence assumption as required by El Karoui and d'Aspremont (2010) and

it explicitly characterizes how the probability distribution affects the final statistical error in finite sample error bounds. Besides, our aggregation algorithm is much simpler than the one in Chen et al. (2016).

We propose a distributed algorithm with only one-shot communication to solve for top K eigenvectors of population covariance matrix $\mathbf{\Sigma}$ when samples are scattered across m servers. We first calculate for each subset of data ℓ its top K eigenvectors $\{\widehat{\mathbf{V}}_K^{(\ell)} = (\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)})\}_{\ell=1}^m$ of the sample covariance matrix there, then derive the average of projection matrices of the corresponding eigenspaces $\widetilde{\mathbf{\Sigma}} = (1/m)\sum_{i=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$, and finally take the top K eigenvectors of $\widetilde{\mathbf{\Sigma}}$ as the final estimator $\widetilde{\mathbf{V}}_K = (\tilde{\mathbf{v}}_1^{(\ell)}, \cdots, \tilde{\mathbf{v}}_K^{(\ell)})$. The communication cost of this method is of order $O(mKd)$. We establish rigorous non-asymptotic analysis of the statistical error $\|\widetilde{\mathbf{V}}_K\widetilde{\mathbf{V}}_K^T - \mathbf{V}_K\mathbf{V}_K^T\|_F$, and show that as long as we have a sufficiently large number of samples in each server, $\widetilde{\mathbf{V}}_K$ enjoys the same statistical error rate as the standard PCA over the full sample. The eigenvalues of $\mathbf{\Sigma}$ are easily estimated once we get good estimators of the eigenvectors, using another round of communication.

The rest of the paper is organized as follows. In Section 2, we introduce the problem setup of the distributed PCA. In Section 3, we elucidate our distributed algorithm for estimating the top K eigenvectors. Section 4 develops the statistical error rates of the aggregated estimator. The results are extended to heterogeneous samples in Section 5. Finally in Section 6 we present extensive simulation results to validate our theories.

2 Problem Setup

We first collect all the notations that will be used. By convention we use regular letters for scalars and bold letters for both matrices and vectors. We denote the set $\{1, 2, 3, ..., d\}$ by $[d]$ for convenience. For two scalars a and b, we write $\min(a, b)$ as $a \wedge b$ and $\max(a, b)$ as $a \vee b$. For two scalar sequences $\{a_n\}_{n\geq 1}$ and $\{b_n\}_{n\geq 1}$, we say $a_n \gtrsim b_n$ ($a_n \precsim b_n$) if there exists a universal constant $C > 0$ such that $a_n \geq Cb_n$ ($a_n \leq Cb_n$). For a random variable $X \in \mathbb{R}$, we define $\|X\|_{\psi_2} = \sup_{p\geq 1}(\mathbb{E}|X|^p)^{\frac{1}{p}}/\sqrt{p}$ and define $\|X\|_{\psi_1} = \sup_{p\geq 1}(\mathbb{E}|X|^p)^{\frac{1}{p}}/p$. Please refer to Vershynin (2010) for equivalent definitions of ψ_2-norm and ψ_1-norm. For two random variables X and Y, we use $X \overset{d}{=} Y$ to denote that X and Y have identical distributions. Define $\mathbf{e}_i$ to be the unit vector whose components are all zero except that

the i-th component equals 1. For $q \geq r$, $\mathcal{O}_{q\times r}$ denotes the space of $q \times r$ matrices with orthonormal columns. For a matrix $\mathbf{A} \in \mathbb{R}^{n\times d}$, we use $\|\mathbf{A}\|_F$, $\|\mathbf{A}\|_N$ and $\|\mathbf{A}\|_{op}$ to denote the Frobenius norm, nuclear norm and operator norm of $\mathbf{A}$ respectively. $\mathrm{Col}(\mathbf{A})$ represents the linear space spanned by column vectors of $\mathbf{A}$. We say $\mathrm{Col}(\mathbf{A}) \perp \mathrm{Col}(\mathbf{B})$ if for any $\mathbf{a} \in \mathrm{Col}(\mathbf{A})$ and $\mathbf{b} \in \mathrm{Col}(\mathbf{B})$, $\mathbf{a}^T\mathbf{b} = 0$. We denote the Moore-Penrose pseudo inverse of a matrix $\mathbf{A} \in \mathbb{R}^{d\times d}$ by $\mathbf{A}^\dagger$.

Suppose we have N i.i.d random samples $\{\mathbf{x}_i\}_{i=1}^N \subseteq \mathbb{R}^d$ with $\mathbb{E}\mathbf{x}_1 = \mathbf{0}$ and covariance matrix $\mathbb{E}(\mathbf{x}_1\mathbf{x}_1^T) = \mathbf{\Sigma}$. By spectral decomposition, $\mathbf{\Sigma} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, where $\mathbf{\Lambda} = \mathrm{diag}(\lambda_1, \lambda_2, \cdots, \lambda_d)$ with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$ and $\mathbf{V} = (\mathbf{v}_1, \cdots, \mathbf{v}_d) \in \mathcal{O}(d)$, the set of all $d \times d$ orthogonal matrices. For a given $K \in [d]$, let $\mathbf{V}_K = (\mathbf{v}_1, \cdots, \mathbf{v}_K)$. Our goal is to estimate $\mathrm{Col}(\mathbf{V}_K)$, i.e., the linear space spanned by the top K eigenvectors of $\mathbf{\Sigma}$. To ensure the identifiability of $\mathrm{Col}(\mathbf{V}_K)$, we assume $\lambda_K > \lambda_{K+1}$.

The standard way of estimating $\mathrm{Col}(\mathbf{V}_K)$ is to use the top K eigenspace of the sample covariance $\widehat{\mathbf{\Sigma}} = \frac{1}{N}\sum_{i=1}^N \mathbf{x}_i\mathbf{x}_i^T$. Let $\widehat{\mathbf{\Sigma}} = \widehat{\mathbf{V}}\widehat{\mathbf{\Lambda}}\widehat{\mathbf{V}}^T$ be spectral decomposition of $\widehat{\mathbf{\Sigma}}$, where $\widehat{\mathbf{\Lambda}} = \mathrm{diag}(\widehat{\lambda}_1, \cdots, \widehat{\lambda}_d)$ with $\widehat{\lambda}_1 \geq \cdots \geq \widehat{\lambda}_d$ and $\widehat{\mathbf{V}} = (\widehat{\mathbf{v}}_1, \cdots, \widehat{\mathbf{v}}_d)$. We use $\mathrm{Col}(\widehat{\mathbf{V}}_K)$, where $\widehat{\mathbf{V}}_K = (\widehat{\mathbf{v}}_1, \cdots, \widehat{\mathbf{v}}_K)$, to estimate $\mathrm{Col}(\mathbf{V}_K)$. To gauge the statistical error, we adopt $\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K) := \|\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T - \mathbf{V}_K\mathbf{V}_K^T\|_F$, which is the Frobenius norm of the difference between projection matrices of two spaces and is a well-defined distance between linear subspaces. In fact, $\rho(\mathbf{V}_K, \widehat{\mathbf{V}}_K)$ is equivalent to the so-called $\sin\Theta$ distance. Denote the singular values of $\widehat{\mathbf{V}}_K^T\mathbf{V}_K$ by $\{\sigma_i\}_{i=1}^K$ in descending order. Recall that $\Theta(\widehat{\mathbf{V}}_K, \mathbf{V}_K) = \mathrm{diag}(\theta_1, \cdots, \theta_K)^T$, the principal angles between $\mathrm{Col}(\mathbf{V}_K)$ and $\mathrm{Col}(\widehat{\mathbf{V}}_K)$, are defined as $(\cos^{-1}\sigma_1, \cdots, \cos^{-1}\sigma_K)^T$. Then we define $\sin\Theta(\widehat{\mathbf{V}}_K, \mathbf{V}_K)$ to be $\mathrm{diag}(\sin\theta_1, \cdots, \sin\theta_K)$, which is a K-by-K matrix. Note that

$$\begin{aligned}\rho^2(\mathbf{V}_K, \widehat{\mathbf{V}}_K) &= \mathrm{Tr}((\mathbf{V}_K\mathbf{V}_K^T - \widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)^2) = \mathrm{Tr}(\mathbf{V}_K\mathbf{V}_K^T) + \mathrm{Tr}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - 2\mathrm{Tr}(\widehat{\mathbf{V}}_K^T\mathbf{V}_K\mathbf{V}_K^T\widehat{\mathbf{V}}_K) \\ &= 2K - \|\widehat{\mathbf{V}}_K^T\mathbf{V}_K\|_F^2 = 2\sum_{i=1}^K(1-\sigma_i^2) = 2\sum_{i=1}^K \sin^2\theta_i = 2\|\sin\Theta(\widehat{\mathbf{V}}_K, \mathbf{V}_K)\|_F^2.\end{aligned} \tag{2.1}$$

Therefore, $\rho(\mathbf{V}_K, \widehat{\mathbf{V}}_K)$ and $\|\sin\Theta(\mathbf{V}_K, \widehat{\mathbf{V}}_K)\|_F$ are equivalent.

Now consider the estimation of top K eigenspace under the distributed data setting, where our $N = m \cdot n$ samples are scattered across m machines with each machine storing n samples*. Application of standard PCA here requires data or covariance aggregation, thus

*Note that here for simplicity we assume the subsample sizes are homogeneous. We can easily extend our

leading to huge communication cost for high-dimensional big data. In addition, for the areas such as genetic, biomedical studies and customer services, it is hard to communicate raw data because of privacy and ownership concerns. To address these problems, we need to avoid naive data aggregation and design a communication-efficient and privacy-preserving distributed algorithm for PCA. In addition, this new algorithm should be statistically accurate in the sense that it enjoys the same statistical error rate as the full sample PCA.

Throughout the paper, we assume that all the random samples $\{\mathbf{x}_i\}_{i=1}^N$ are i.i.d sub-Gaussian. We adopt the definition of sub-Gaussian random vectors in Koltchinskii and Lounici (2017) and Reiss and Wahl (2016) as specified below. It is not hard to show that the following definition is equivalent to the definition $\|(\mathbf{\Sigma}^{1/2})^\dagger X\|_{\psi_2} \leq M$ used in Vershynin (2010), Wang and Fan (2017), and many other authors.

Definition 2.1. *We say the random vector* $\mathbf{x} \in \mathbb{R}^d$ *is sub-Gaussian if there exists* $M > 0$ *such that* $\|\mathbf{u}^T\mathbf{x}\|_{\psi_2} \leq M\sqrt{\mathbb{E}(\mathbf{u}^T\mathbf{x})^2}$, $\forall \mathbf{u} \in \mathbb{R}^d$.

We emphasize here that the global i.i.d assumption on $\{\mathbf{x}_i\}_{i=1}^N$ can be further relaxed. In fact, our statistical analysis only requires the following three conditions: (i) within each server ℓ, data are i.i.d. ; (ii) across different servers, data are independent; (iii) the covariance matrices of the data in each server $\{\mathbf{\Sigma}^{(\ell)}\}_{\ell=1}^m$ share similar top K eigenspaces. We will further study this heterogeneous regime in Section 5. To avoid future confusion, unless we clarify that we are under this covariance heterogeneity regime, we always assume i.i.d. data across servers.

3 Methodology

We now introduce our distributed PCA algorithm. For $\ell \in [m]$, let $\{\mathbf{x}_i^{(\ell)}\}_{i=1}^n$ denote the samples stored on the ℓ-th machine. We specify the steps of the algorithm in Algorithm 1.

In other words, each server first calculates the top K eigenvectors of the local sample covariance matrix, and then transmits these eigenvectors $\{\widehat{\mathbf{V}}_K^{(\ell)}\}_{\ell=1}^m$ to a central server. It is worth noting that in this algorithm, we do not really need to compute $\{\widehat{\mathbf{\Sigma}}^{(\ell)}\}_{\ell=1}^m$ and $\widetilde{\mathbf{\Sigma}}$. $\{\widehat{\mathbf{V}}_K^{(\ell)}\}_{\ell=1}^m$ and $\widetilde{\mathbf{V}}_K$ can be derived from top-K SVD of data matrices. This is far more expeditious than the entire SVD and highly scalable, by using, for example, the power

analysis to the case of heterogeneous sub-sample sizes with similar theoretical results.

Algorithm 1 Distributed PCA

1. On each server, compute locally the K leading eigenvectors $\widehat{\mathbf{V}}_K^{(\ell)} = (\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)}) \in \mathbb{R}^{d\times K}$ of the sample covariance matrix $\widehat{\mathbf{\Sigma}}^{(\ell)} = (1/n)\sum_{i=1}^n \mathbf{x}_i^{(\ell)}\mathbf{x}_i^{(\ell)^T}$. Send $\widehat{\mathbf{V}}_K^{(\ell)}$ to a central processor.

2. On the central processor, compute $\widetilde{\mathbf{\Sigma}} = (1/m)\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$, and its K leading eigenvectors $\{\tilde{\mathbf{v}}_j\}_{j=1}^K$. Output: $\widetilde{\mathbf{V}}_K = (\widetilde{\mathbf{v}}_1, \cdots, \widetilde{\mathbf{v}}_K) \in \mathbb{R}^{d\times K}$.

method (Golub and Van Loan, 2012). As regard to the estimation of the top eigenvalues of $\mathbf{\Sigma}$, we can send the aggregated eigenvectors $\{\tilde{\mathbf{v}}_j\}_{j=1}^K$ back to the m servers, where each one computes $\{\lambda_j^{(\ell)}\}_{j=1}^K = \{\tilde{\mathbf{v}}_j^T\widehat{\mathbf{\Sigma}}^{(\ell)}\tilde{\mathbf{v}}_j\}_{j=1}^K$. Then the central server collect all the eigenvalues and deliver the average eigenvalues $\{\tilde{\lambda}_j\}_{j=1}^K = \{\frac{1}{m}\sum_{\ell=1}^m \lambda_j^{(\ell)}\}_{j=1}^K$ as the estimators of all eigenvalues.

As we can see, the communication cost of the proposed distributed PCA algorithm is of order $O(mKd)$. In contrast, to share all the data or entire covariance, the communication cost will be of order $O(md\min(n,d))$. Since in most cases $K = o(\min(n,d))$, our distributed PCA requires much less communication cost than naive data aggregation.

4 Statistical Error Analysis

Algorithm 1 delivers $\widetilde{\mathbf{V}}_K$ to estimate the top K eigenspace of $\mathbf{\Sigma}$. In this section we analyze the statistical error of $\widetilde{\mathbf{V}}_K$, i.e., $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$. The main message is that $\widetilde{\mathbf{V}}_K$ enjoys the same statistical error rate as the full sample counterpart $\widehat{\mathbf{V}}_K$ as long as the subsample size n is sufficiently large.

We first conduct a bias and variance decomposition of $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$, which serves as the key step in establishing our theoretical results. Recall that $\widetilde{\mathbf{\Sigma}} = (1/m)\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$ and $\widetilde{\mathbf{V}}_K$ consists of the top K eigenvectors of $\widetilde{\mathbf{\Sigma}}$. Define $\mathbf{\Sigma}^* := \mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})$ and denote its top K eigenvectors by $\mathbf{V}_K^* = (\mathbf{v}_1^*, \cdots, \mathbf{v}_K^*) \in \mathbb{R}^{d\times K}$. When the number of machines is large, $\widetilde{\mathbf{\Sigma}}$ will converge to $\mathbf{\Sigma}^*$, and naturally we expect $\mathrm{Col}(\widetilde{\mathbf{V}}_K)$ to converge to $\mathrm{Col}(\mathbf{V}_K^*)$ as well. This line of thinking inspires us to decompose the statistical error $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ into the following

bias and sample variance terms:

$$\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K) \leq \underbrace{\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)}_{\text{sample variance term}} + \underbrace{\rho(\mathbf{V}_K^*, \mathbf{V}_K)}_{\text{bias term}}. \tag{4.1}$$

The first term is stochastic and the second term is deterministic. Here we elucidate on why we call $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$ the sample variance term and $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ the bias term respectively.

1. Sample variance term $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$:

   By Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)) and (2.1), we have

   $$\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*) \leq \frac{2\sqrt{2}\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F}{\lambda_K(\mathbf{\Sigma}^*) - \lambda_{K+1}(\mathbf{\Sigma}^*)}. \tag{4.2}$$

   As we can see, $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$ depends on how the average $\widetilde{\mathbf{\Sigma}} = (1/m)\sum_{\ell=1}^{m} \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$ concentrates to its mean $\mathbf{\Sigma}^*$. This explains why we call $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$ the sample variance term. We will show in the sequel that for sub-Gaussian random samples, $\{\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F\}_{\ell=1}^m$ and $\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F$ are sub-exponential random variables and under appropriate regularity assumptions it holds that

   $$\|\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F\|_{\psi_1} \leq \frac{C}{\sqrt{m}}\|\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)^T} - \mathbf{\Sigma}^*\|_F\|_{\psi_1} \tag{4.3}$$

   for some universal constant $C > 0$. If we regard ψ_1-norm as standard deviation proxy, this result is a counterpart to the formula for the standard deviation of the sample mean under the context of matrix concentration. By (4.3), the average of projection matrices $\widetilde{\mathbf{\Sigma}}$ enjoys a similar square-root convergence.

2. Bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$:

   The error $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ is deterministic and independent of how many machines we have, and is therefore called the bias term. We will show this bias term is exactly zero when the random sample has a symmetric innovation (to be defined later). In general, we will show that the bias term is negligible in comparison with the sample variance term when the number of nodes m is not unreasonably large.

In the following subsections, we will analyze the sample variance term and bias term respectively and then combine these results to obtain the convergence rate for $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$.

4.1 Analysis of the Sample Variance Term

To analyze $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$, as shown by (4.2), we need to derive the order of the numerator $\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F$ and denominator $\lambda_K(\mathbf{\Sigma}^*) - \lambda_{K+1}(\mathbf{\Sigma}^*)$. We first focus on the matrix concentration term $\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F = \|(1/m)\sum_{\ell=1}^{m}\Big(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\Big)\|_F$. Note that $\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*$ is an average of m centered random matrices. To establish the correspondent concentration inequality, we first investigate each individual term in the average, i.e., $\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*$ for $\ell \in [m]$. In the following lemma, we show that when random samples are sub-Gaussian, $\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F$ is sub-exponential and we can give an explicit upper bound of its ψ_1-norm.

Lemma 1. *Suppose in the ℓ-th server, we have n i.i.d. sub-Gaussian random samples $\{\mathbf{x}_i\}_{i=1}^n$ in $\mathbb{R}^d$ with covariance matrix $\mathbf{\Sigma}$. Then, there exists a constant $C > 0$ such that, for any $t > 1$, with probability at least $1 - \exp(-t)$,*

$$\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F \leq \frac{C\sqrt{K}\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\Big(\sqrt{\frac{r(\mathbf{\Sigma})}{n}} \vee \frac{r(\mathbf{\Sigma})}{n} \vee \sqrt{\frac{t}{n}} \vee \frac{t}{n}\Big),$$

where $r(\mathbf{\Sigma}) := (\mathbb{E}\|\mathbf{x}\|_2)^2/\|\mathbf{\Sigma}\|_{op}$, which is smaller than the so-called effective rank of $\mathbf{\Sigma}$, i.e., $\mathrm{Tr}(\mathbf{\Sigma})/\|\mathbf{\Sigma}\|_{op}$. In other words, $\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F$ is a sub-exponential random variable and satisfies

$$\|\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F\|_{\psi_1} \leq \frac{C\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\sqrt{\frac{Kr(\mathbf{\Sigma})}{n}}.$$

Note that here we use the Frobenius norm to measure the distance between two matrices. Therefore, it is equivalent to treat $\{\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}\}_{\ell=1}^K$ and $\mathbf{\Sigma}^*$ as d^2-dimensional vectors and apply the concentration inequality for random vectors to bound $\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F$. As we will demonstrate in the proof of Theorem 1, $\|\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F\|_{\psi_1} = 1/\sqrt{m}\|\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F\|_{\psi_1}$.

With regard to $\lambda_K(\mathbf{\Sigma}^*) - \lambda_{K+1}(\mathbf{\Sigma}^*)$, when the individual node has enough samples, $\widehat{\mathbf{V}}_K^{(\ell)}$ and $\mathbf{V}_K$ will be close to each other and so will be $\mathbf{\Sigma}^* = \mathbb{E}(\mathbf{V}_K^{(\ell)}\mathbf{V}_K^{(\ell)^T})$ and $\mathbf{V}_K\mathbf{V}_K^T$. Given that $\lambda_K(\mathbf{V}_K\mathbf{V}_K^T) = 1$ and $\lambda_{K+1}(\mathbf{V}_K\mathbf{V}_K^T) = 0$, we accordingly expect $\lambda_K(\mathbf{\Sigma}^*)$ and $\lambda_{K+1}(\mathbf{\Sigma}^*)$ be separated by a positive constant as well.

All the arguments above lead to the following theorem on $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$.

Theorem 1. *Suppose* $\mathbf{x}, \mathbf{x}_1, \cdots, \mathbf{x}_N$ *are i.i.d. sub-Gaussian random vectors in* $\mathbb{R}^d$ *with covariance matrix* $\mathbf{\Sigma}$ *and they are scattered across* m *machines. If* $\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \leq 1/4$ *for all* $\ell \in [m]$*, then we have*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)\|_{\psi_1} \leq \frac{C\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}} \sqrt{\frac{Kr(\mathbf{\Sigma})}{N}},$$

where $N = mn$ *is the total sample size and* C *is some universal constant.*

4.2 Analysis of the Bias Term

In this section, we study the bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ in (4.1). We first focus on a special case where the bias term is exactly zero. For a random vector $\mathbf{x}$ with covariance $\mathbf{\Sigma} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, let $\mathbf{z} = \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{V}^T\mathbf{x}$. We say $\mathbf{x}$ has symmetric innovations if $\mathbf{z} \overset{d}{=} (\mathbf{I}_d - 2\mathbf{e}_j\mathbf{e}_j^T)\mathbf{z}, \quad \forall j \in [d]$. In other words, flipping the sign of one component of $\mathbf{z}$ will not change the distribution of $\mathbf{z}$. Note that if $\mathbf{z}$ has density, this is equivalent to say that its density function has the form $p(z_1^2, z_2^2, \cdots, z_d^2)$. All elliptical distributions centered at the origin belong to this family. In addition, if $\mathbf{z}$ has symmetric and independent entries, $\mathbf{x}$ has also symmetric innovations. It turns out that when the random samples have symmetric innovations, $\mathbf{\Sigma}^* := \mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})$ and $\mathbf{\Sigma}$ share exactly the same set of eigenvectors. When we were finishing the paper, we noticed that Chen et al. (2016) have independently got a similar result for the Gaussian case.

Theorem 2. *Suppose in the* ℓ*-th server we have* n *i.i.d. random samples* $\{\mathbf{x}_i\}_{i=1}^n$ *with covariance* $\mathbf{\Sigma}$*. If* $\{\mathbf{x}_i\}_{i=1}^n$ *are symmetric, then* $\mathbf{V}^T\mathbf{\Sigma}^*\mathbf{V}$ *is diagonal, i.e.,* $\mathbf{\Sigma}^*$ *and* $\mathbf{\Sigma}$ *share the same set of eigenvectors. Furthermore, if* $\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \leq 1/4$*, then* $\rho(\mathbf{V}_K^*, \mathbf{V}_K) = 0$*.*

Here we explain why we need the condition $\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \leq 1/4$ to achieve zero bias. First of all, the condition is similar to a bound on the "variance" of the random matrix $\widehat{\mathbf{V}}_K^{(\ell)}$ whose covariance $\mathbf{\Sigma}^*$ is under investigation. As demonstrated above, with the symmetric innovations, $\mathbf{\Sigma}^*$ has the same set of eigenvectors as $\mathbf{\Sigma}$, but we still cannot guarantee that the top K eigenvectors of $\mathbf{\Sigma}^*$ match with those of $\mathbf{\Sigma}$. For example, the $(K+1)$th top eigenvector of $\mathbf{\Sigma}$ might be the Kth top eigenvector of $\mathbf{\Sigma}^*$. In order to ensure the top K eigenspace of $\mathbf{\Sigma}^*$ is exactly the same as that of $\mathbf{\Sigma}$, we require $\widehat{\mathbf{V}}_K^{(\ell)}$ can not deviate too far from $\mathbf{V}_K$ so that $\mathbf{\Sigma}^*$ is close enough to $\mathbf{V}_K\mathbf{V}_K^T$.

For general distributions, the bias term is not necessarily zero. However, we discover that when the subsample size is large enough, the bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ is of high-order compared with the statistical error of $\widehat{\mathbf{V}}_K^{(\ell)}$ on the individual subsample. By the decomposition (4.1) and Theorem 1, we can therefore expect the aggregated estimator $\widetilde{\mathbf{V}}_K$ to enjoy sharper statistical error rate than PCA on the individual subsample. In other words, the aggregation does improve the statistical efficiency. A similar phenomenon also appears in statistical error analysis of the average of the debiased Lasso estimators in Battey et al. (2015) and Lee et al. (2015). Recall that in sparse linear regression, the Lasso estimator $\widehat{\boldsymbol{\beta}}$ satisfies that $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 = O_P(\sqrt{s \log d/n})$, where $\boldsymbol{\beta}^*$ is the true regression vector, s is the number of nonzero coefficients of $\boldsymbol{\beta}^*$ and d is the dimension. The debiasing step reduces the bias of $\widehat{\boldsymbol{\beta}}$ to the order $O_P(s \log d/n)$, which is negligible, compared with the statistical error of $\widehat{\boldsymbol{\beta}}$ and thus enables the averaging of the debiased Lasso estimators to enhance the statistical efficiency.

Below we present a second-order Davis-Kahan theorem that explicitly characterizes the first and second order errors on top K eigenspace due to matrix perturbation. This is a genuine generalization of the former high-order perturbation theories on a single eigenvector, e.g., Lemma 1 in Kneip and Utikal (2001) and Theorem 2 in El Karoui and d'Aspremont (2010). By nature it is deterministic and serves as the core of the bias analysis in our work. Let us first introduce the notation.

Let $\mathbf{A}, \widehat{\mathbf{A}} \in \mathbb{R}^{d\times d}$ be symmetric matrices with eigenvalues $\lambda_1 \geq \cdots \geq \lambda_d$, and $\widehat{\lambda}_1 \geq \cdots \geq \widehat{\lambda}_d$, respectively. Let $\{\mathbf{u}_j\}_{j=1}^d$, $\{\widehat{\mathbf{u}}_j\}_{j=1}^d$ be two orthonormal bases of $\mathbb{R}^d$ such that $\mathbf{A}\mathbf{u}_j = \lambda_j \mathbf{u}_j$ and $\widehat{\mathbf{A}}\widehat{\mathbf{u}}_j = \widehat{\lambda}_j \widehat{\mathbf{u}}_j$ for all $j \in [d]$. Fix $s \in \{0, 1, \cdots, d-K\}$ and assume that $\Delta = \min\{\lambda_s - \lambda_{s+1}, \lambda_{s+K} - \lambda_{s+K+1}\} > 0$, where $\lambda_0 = +\infty$ and $\lambda_{d+1} = -\infty$. Define $\mathbf{U} = (\mathbf{u}_{s+1}, \cdots, \mathbf{u}_{s+K})$, $\widehat{\mathbf{U}} = (\widehat{\mathbf{u}}_{s+1}, \cdots, \widehat{\mathbf{u}}_{s+K})$ and $\mathbf{H} = \widehat{\mathbf{U}}^T\mathbf{U}$. Let $\mathbf{H} = \mathbf{X}\mathbf{D}\mathbf{Y}^T$ be its singular value decomposition and $\widehat{\mathbf{H}} = \mathbf{X}\mathbf{Y}^T$. Define $\mathbf{E} = \widehat{\mathbf{A}} - \mathbf{A}$, $S = \{s+1, \cdots, s+K\}$, $\mathbf{G}_j = \sum_{i\notin S}(\lambda_i - \lambda_{s+j})^{-1}\mathbf{u}_i\mathbf{u}_i^T$ for $j \in [K]$, and

$$f : \mathbb{R}^{d\times K} \to \mathbb{R}^{d\times K}, \quad (\mathbf{v}_1, \cdots, \mathbf{v}_K) \mapsto (-\mathbf{G}_1\mathbf{v}_1, \cdots, -\mathbf{G}_K\mathbf{v}_K).$$

Lemma 2. *Assuming that* $\varepsilon = \sqrt{8K}\|\mathbf{E}\|_{op}/\Delta \leq 1/\sqrt{2}$*, we have*

$$\begin{aligned}
&\min_{\mathbf{O}\in\mathcal{O}_{K\times K}} \|\widehat{\mathbf{U}}\mathbf{O} - \mathbf{U}\|_F = \|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F \leq \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T\|_F \leq \varepsilon, \\
&\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{E}\mathbf{U})\|_F \leq 5\varepsilon^2, \\
&\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - [\mathbf{U}\mathbf{U}^T + f(\mathbf{E}\mathbf{U})\mathbf{U}^T + \mathbf{U}f(\mathbf{E}\mathbf{U})^T]\|_F \leq 12\varepsilon^2.
\end{aligned} \tag{4.4}$$

The first bound is just the standard Davis-Kahan inequality, see Theorem 2 in Yu et al. (2015). The next two investigate the difference between column space spanned by $\widehat{\mathbf{U}}$ and $\mathbf{U}$. As in regular Taylor expansion, we decompose the difference into the leading linear terms and quadratic residuals with respect to the perturbation. Now we apply Lemma 2 to the context of principal eigenspace estimation. Let $\mathbf{A} = \mathbf{\Sigma}$, $\mathbf{E} = \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma}$ and $S = [K]$. It thus follows that $\mathbf{U} = \mathbf{V}_K$ and $\widehat{\mathbf{U}} = \widehat{\mathbf{V}}_K$. Since $\mathbb{E}\mathbf{E} = \mathbf{0}$, the expectation of the linear leading term $\mathbb{E}(f(\mathbf{E}\mathbf{V}_K)\mathbf{V}_K^T + \mathbf{V}_K f(\mathbf{E}\mathbf{V}_K)^T)$ is then zero. Based on this fact, from the last inequality in (4.4) we can derive that the bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ will be a high-order term compared with the linear leading term. More rigorously, by Davis-Kahan's Theorem, we can control the bias as follows:

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq 2\sqrt{2} \cdot \|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F = 2\sqrt{2} \cdot \|\mathbb{E}[\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T - \mathbf{V}_K\mathbf{V}_K^T]\|_F. \tag{4.5}$$

By using the fact that $\mathbb{E}\mathbf{E} = 0$ and f is linear, we have

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq 2\sqrt{2} \cdot \|\mathbb{E}[\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T - (\mathbf{V}_K\mathbf{V}_K^T + f(\mathbf{E}\mathbf{V}_K)\mathbf{V}_K^T + \mathbf{V}_K f(\mathbf{E}\mathbf{V}_K)^T)]\|_F$$

By Jensen's inequality, the above is further bounded by

$$2\sqrt{2} \cdot \mathbb{E}\|\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T - (\mathbf{V}_K\mathbf{V}_K^T + f(\mathbf{E}\mathbf{V}_K)\mathbf{V}_K^T + \mathbf{V}_K f(\mathbf{E}\mathbf{V}_K)^T)\|_F$$

Finally, by Lemma 2, we have

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq 24\sqrt{2}\mathbb{E}(\varepsilon^2). \tag{4.6}$$

Using the above arguments, we obtain the following theorem.

Theorem 3. *Suppose we have* n *i.i.d. sub-Gaussian random samples* $\{\mathbf{x}_i\}_{i=1}^n$ *in the* ℓ*-th*

server with covariance $\mathbf{\Sigma}$*. Then, there exists some constant* $C_2 > 0$ *such that*

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq \frac{C_2 K \|\mathbf{\Sigma}\|_{op}^2 \cdot r(\mathbf{\Sigma})}{n(\lambda_K - \lambda_{K+1})^2}.$$

4.3 Properties of distributed PCA

We now combine the results we obtained in the previous two subsections to derive the statistical error rate of $\widetilde{\mathbf{V}}_K$. We first present a theorem under the setting of global i.i.d. data and discuss its optimality.

Theorem 4. *Suppose we have* N *i.i.d. sub-Gaussian random samples with covariance* $\mathbf{\Sigma}$*. They are scattered across* m *servers, each of which stores* n *samples.*

1. *Sign symmetric innovation: There exist some positive constants* C_1 *and* η_1 *such that*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq \frac{C_1 \|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}} \sqrt{\frac{Kr(\mathbf{\Sigma})}{N}} \tag{4.7}$$

   so long as $n \geq \eta_1 Kr(\mathbf{\Sigma})(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}))^2$*.*

2. *General distribution: There exist some positive constants* C_2, C_3 *and* η_2 *such that*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq \frac{C_2 \|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}} \sqrt{\frac{Kr(\mathbf{\Sigma})}{N}} + \frac{C_3 \|\mathbf{\Sigma}\|_{op}^2 \cdot Kr(\mathbf{\Sigma})}{n(\lambda_K - \lambda_{K+1})^2} \tag{4.8}$$

   provided that $n \geq \eta_2 Kr(\mathbf{\Sigma})(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}))^2$*. In addition, there exist positive constants* C_4 *and* γ *such that if we further assume* $m \leq \gamma n(\lambda_K - \lambda_{K+1})^2/(Kr(\mathbf{\Sigma})\|\mathbf{\Sigma}\|_{op}^2)$*,*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq \frac{C_4 \|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}} \sqrt{\frac{Kr(\mathbf{\Sigma})}{N}}. \tag{4.9}$$

As we can see, with appropriate scaling conditions on n, m, d, $\widetilde{\mathbf{V}}_K$ can achieve the statistical error rate of order $O_P(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}) \cdot \sqrt{Kr(\mathbf{\Sigma})/N})$. The result is applicable to the whole sample or traditional PCA, in which $m = 1$. This demonstrates that the distributed PCA admits the same rate as the traditional PCA, as long as the technical conditions are satisfied.

We now illustrate this result through a simple spiked covariance model introduced by Johnstone (2001). Assume

$$\mathbf{\Lambda} = \mathrm{diag}(\lambda, \underbrace{1, \cdots, 1}_{d-1}), \tag{4.10}$$

where $\lambda > 1$. Since $r(\mathbf{\Sigma}) \leq \mathrm{Tr}(\mathbf{\Sigma})/\|\mathbf{\Sigma}\|_{op} = (\lambda + d - 1)/\lambda$, when $\lambda = O(d)$ [including $\lambda = o(d)$], it is easy to show from (4.7) or (4.9) that

$$\|\rho(\widetilde{\mathbf{V}}_1, \mathbf{V}_1)\|_{\psi_1} = O\Big(\sqrt{\frac{d}{N\lambda}}\Big). \tag{4.11}$$

We now compare this rate with the previous results under the spiked model (4.9). In Paul and Johnstone (2012), the authors derived the ℓ_2 risk of the empirical eigenvectors when random samples are Gaussian. It is not hard to derive from Theorem 1 therein that given N i.i.d d-dimensional Gaussian samples, when N, d, λ go to infinity,

$$\mathbb{E}\|\widehat{\mathbf{V}}_1 - \mathbf{V}_1\|_2^2 = O\Big(\frac{d}{N\lambda}\Big),$$

where we require $\widehat{\mathbf{V}}_1^T \mathbf{V}_1 > 0$, i.e., we always adjust the sign of $\widehat{\mathbf{V}}_1$ such that the angle between $\widehat{\mathbf{V}}_1$ and $\mathbf{V}_1$ is acute. Note that

$$\begin{aligned}\rho^2(\widetilde{\mathbf{V}}_1, \mathbf{V}_1) &= \|\widetilde{\mathbf{V}}_1\widetilde{\mathbf{V}}_1^T - \mathbf{V}_1\mathbf{V}_1^T\|_F^2 = 2(1 - \widetilde{\mathbf{V}}_1^T\mathbf{V}_1)(1 + \widetilde{\mathbf{V}}_1^T\mathbf{V}_1) \\ &\geq 2(1 - \widetilde{\mathbf{V}}_1^T\mathbf{V}_1) = \|\widetilde{\mathbf{V}}_1 - \mathbf{V}_1\|_2^2.\end{aligned}$$

Hence from (4.11), we obtain $\mathbb{E}\|\widetilde{\mathbf{V}}_1 - \mathbf{V}_1\|_2^2 = O(d/(N\lambda))$, which has the same rate as $\widehat{\mathbf{V}}_1$. The aggregated estimator $\widetilde{\mathbf{V}}_1$ performs as well as the full sample estimator $\widehat{\mathbf{V}}_1$ in terms of the mean square error. In fact, by taking $m = 1$ (the whole sample), the result of Paul and Johnstone (2012) is a corollary of ours, as demonstrated above.

In addition, our result is consistent with the minimax lower bound developed in Cai et al. (2013). For $\kappa, \lambda > 0$, define

$$\Theta(\kappa, \lambda) := \{\mathbf{\Sigma} \succeq 0 : \lambda + 1 \leq \lambda_K \leq \lambda_1 \leq \kappa\lambda + 1, \lambda_j = 1 \text{ for } j \geq K\}.$$

Theorem 8 in Cai et al. (2013) shows that under the Gaussian distribution with $\mathbf{\Sigma} \in \Theta(\kappa, \lambda)$,

the minimax lower bound of $\mathbb{E}\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K)^2$ is

$$\inf_{\widehat{\mathbf{V}}} \sup_{\mathbf{\Sigma}\in\Theta(\kappa,\lambda)} \mathbb{E}\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K)^2 \geq C \min\Big(K, (d-K), \frac{K(\lambda+1)(d-K)}{N\lambda^2}\Big),$$

for some universal constant $C > 0$. When $d > K/2$ and λ grows to infinity, the RHS of the inequality above is reduced to $\min\Big(K, Kd/(N\lambda)\Big)$. Given that $\|\mathbf{\Sigma}\|_{op} \leq \kappa\lambda + 1$, $\lambda_K - \lambda_{K+1} \geq \lambda$ and $r(\mathbf{\Sigma}) \leq C'(K\lambda + d)/\lambda$ for some universal $C' > 0$, our upper bound in (4.7) matches this lower bound when $\lambda = O(d)$.

Notice that in Theorem 4, the prerequisite for $\widetilde{\mathbf{V}}_K$ to enjoy the sharp statistical error rate is a lower bound on the subsample size n, i.e.,

$$n \gtrsim Kr(\mathbf{\Sigma})(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}))^2. \tag{4.12}$$

It is natural to ask whether this required lower bound is sharp, or in other words, is it possible for $\widetilde{\mathbf{V}}_K$ to achieve the same statistical error rate with a smaller sample size on each machine? The answer is no. The following theorem presents a distribution family under which $\mathrm{Col}(\widetilde{\mathbf{V}}_K)$ is even perpendicular to $\mathrm{Col}(\mathbf{V}_K)$ with high probability when n is smaller than the threshold given in (4.12). This means that having a smaller sample size on each machine is too uninformative and even malicious such that the aggregation step completely fails in improving estimation consistency.

Theorem 5. *Consider a Bernoulli random variable W with $P(W = 0) = P(W = 1) = 1/2$, a Rademacher random variable $P(Y = 1) = P(Y = -1) = 1/2$, and a random vector $\mathbf{Z} \in \mathbb{R}^d$ that is uniformly distributed over the d-dimensional unit sphere. We say a random vector $\mathbf{X} \in \mathbb{R}^d$ follows the distribution $\mathcal{D}(\lambda)$ if*

$$\mathbf{X} \overset{d}{=} \begin{pmatrix} \mathbf{1}_{\{W=0\}}\sqrt{2\lambda}Y \\ \mathbf{1}_{\{W=1\}}\sqrt{2(d-1)}\mathbf{Z} \end{pmatrix}.$$

Now suppose we have $\{\mathbf{X}_i\}_{i=1}^N$ as N i.i.d. random samples of $\mathbf{X}$. They are stored across m

servers, each of which has n samples. When $32 \log d \le n \le (d-1)/(3\lambda)$, we have

$$P(\widetilde{\mathbf{V}}_1 \perp \mathbf{V}_1) \ge \begin{cases} 1 - d^{-1}, & \text{if } m \le d^3, \\ 1 - e^{-d/2}, & \text{if } m > d^3. \end{cases}$$

It is easy to verify that $\mathcal{D}(\lambda)$ is symmetric, sub-Gaussian and satisfies $\mathbb{E}\mathbf{X} = \mathbf{0}$ and $\mathbb{E}\mathbf{X}\mathbf{X}^T = \mathrm{diag}(\lambda, 1, \cdots, 1)$. According to (4.12), when $\lambda = O(d)$, we require $n \gtrsim d/\lambda$ to achieve the rate as demonstrated in (4.7). Theorem 5 shows that if we have fewer samples than this threshold, the aggregated estimator $\widetilde{\mathbf{V}}_1$ will be perpendicular to the true top eigenvector $\mathbf{V}_1$ with high probability. Therefore, our lower bound on the subsample size n is sharp.

5 Extension to Heterogeneous Samples

We now relax global *i.i.d.* assumptions in the previous section to the setting of heterogeneous covariance structures across servers. Suppose data on the server ℓ has covariance matrix $\mathbf{\Sigma}^{(\ell)}$, whose top K eigenvalues and eigenvectors are denoted by $\{\lambda_k^{(\ell)}\}_{k=1}^K$ and $\mathbf{V}_K^{(\ell)} = (\mathbf{v}_1^{(\ell)}, \cdots, \mathbf{v}_K^{(\ell)})$ respectively. We will study two specific cases of heterogeneous covariances: one requires all covariances to share exactly the same principal eigenspaces, while the other considers the heterogeneous factor models with common factor eigen-structures.

5.1 Common Principal Eigenspaces

We assume $\{\mathbf{\Sigma}^{(\ell)}\}_{\ell=1}^m$ share the same top K eigenspace, i.e., there exists some orthonormal matrix $\mathbf{V}_K \in \mathbb{R}^{d\times K}$ such that $\mathbf{V}_K^{(\ell_1)}\mathbf{V}_K^{(\ell_1)^T} = \mathbf{V}_K^{(\ell_2)}\mathbf{V}_K^{(\ell_2)^T} = \mathbf{V}_K\mathbf{V}_K^T$ for all $\ell_1 \neq \ell_2$. With slight abuse of notations, here we redefine $\mathbf{\Sigma}^* := \frac{1}{m}\sum_{\ell=1}^m \mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$ and $\mathbf{V}_K^* \in \mathbb{R}^{d\times K}$ to be the top K eigenvectors of $\mathbf{\Sigma}^*$. The following lemma shows that $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)$, the sample variance term, enjoys a similar rate as in Theorem 1, and is an extension of Theorem 1.

Lemma 3. *Under the heterogeneous covariance setting, if $\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \le 1/4$ for all*

$\ell \in [m]$, *then there exists a universal constant* $C > 0$ *such that*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)\|_{\psi_1} \leq \sqrt{\frac{1}{m}\sum_{\ell=1}^{m} \frac{CKr(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2}{N(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)})^2}}.$$

We now consider the bias. When the data are driven by symmetric innovations, if $\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \leq 1/4$ for all $\ell \in [m]$, then Theorem 2 implies that the top K eigenspace of $\mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$ is $\mathrm{Col}(\mathbf{V}_K)$. Therefore, the top K eigenspace of $\mathbf{\Sigma}^*$ is still $\mathrm{Col}(\mathbf{V}_K)$ and $\rho(\mathbf{V}_K, \mathbf{V}_K^*) = 0$. For general distributions, again as derived in (4.5), we have

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq 2\sqrt{2} \cdot \|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F \leq \frac{2\sqrt{2}}{m}\sum_{\ell=1}^{m}\|\mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{V}_K\mathbf{V}_K^T\|_F.$$

Following the steps to reach the bound in (4.6), we can similarly characterize the bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ for general distribution of $\{\mathbf{x}_i\}_{i=1}^N$ under the heterogeneous regime. This together with Lemma 3 yields the following result.

Theorem 6. *Suppose we have in total* N *sub-Gaussian data. They are scattered across* m *servers, each of which stores* n *i.i.d. data with covariance* $\mathbf{\Sigma}^{(\ell)}$. $\{\mathbf{\Sigma}^{(\ell)}\}_{\ell=1}^m$ *share the same top* K *eigenspace, i.e.,* $\mathbf{V}_K^{(\ell_1)}\mathbf{V}_K^{(\ell_1)^T} = \mathbf{V}_K^{(\ell_2)}\mathbf{V}_K^{(\ell_2)^T} = \mathbf{V}_K\mathbf{V}_K^T$ *for* $\ell_1 \neq \ell_2$ *and some* $\mathbf{V}_K \in \mathcal{O}_{d\times K}$. *For each* $\ell \in [m]$, *let*

$$S^{(\ell)} := \frac{\|\mathbf{\Sigma}^{(\ell)}\|_{op}}{\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)}}\sqrt{\frac{Kr(\mathbf{\Sigma}^{(\ell)})}{N}} \quad \textit{and} \quad B^{(\ell)} := \frac{K\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2 \cdot r(\mathbf{\Sigma}^{(\ell)})}{n(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)})^2}.$$

1. *Sign symmetric innovation: There exist some positive constants* C_1 *and* η_1 *such that*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq C_1\sqrt{\frac{1}{m}\sum_{\ell=1}^{m} S^{(\ell)^2}} \tag{5.1}$$

so long as $n \geq \eta_1 \max_{\ell\in[m]} Kr(\mathbf{\Sigma})\|\mathbf{\Sigma}\|_{op}^2/(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)})^2$.

2. *General distribution: There exist some positive constants C_2, C_3 and η_2 such that*

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq C_2 \sqrt{\frac{1}{m}\sum_{\ell=1}^{m} S^{(\ell)^2}} + \frac{C_3}{m}\sum_{\ell=1}^{m} B^{(\ell)} \tag{5.2}$$

provided that $n \geq \eta_2 \max_{\ell\in[m]} Kr(\mathbf{\Sigma})\|\mathbf{\Sigma}\|_{op}^2/(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)})^2$. In addition, there exist positive constants C_4 and γ such that

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq C_4 \sqrt{\frac{1}{m}\sum_{\ell=1}^{m} S^{(\ell)^2}} \tag{5.3}$$

as long as if we have $\sum_{\ell=1}^{m} B^{(\ell)} \leq \gamma$.

5.2 Heterogeneous Factor Models

Suppose on the server ℓ, the data conform to a factor model as below.

$$\mathbf{x}_i^{(\ell)} = \mathbf{B}^{(\ell)}\mathbf{f}_i^{(\ell)} + \mathbf{u}_i^{(\ell)}, \quad i \in [n],$$

where $\mathbf{B}^{(\ell)} \in \mathbb{R}^{d\times K}$ is the loading matrix, $\mathbf{f}_i^{(\ell)} \in \mathbb{R}^K$ is the factor that satisfies $\mathrm{Cov}(\mathbf{f}_i^{(\ell)}) = \mathbf{I}$ and $\mathbf{u}_i^{(\ell)} \in \mathbb{R}^d$ is the residual vector. It is not hard to see that $\mathbf{\Sigma}^{(\ell)} = \mathrm{Cov}(\mathbf{x}_i^{(\ell)}) = \mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)^T} + \mathbf{\Sigma}_u^{(\ell)}$, where $\mathbf{\Sigma}_u^{(\ell)}$ is the covariance matrix of $\mathbf{u}_i^{(\ell)}$.

Let $\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)^T} = \mathbf{V}_K^{(\ell)}\mathbf{\Lambda}_K^{(\ell)}\mathbf{V}_K^{(\ell)^T}$ be the spectral decomposition of $\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)^T}$. We assume that there exists a projection matrix $\mathbf{P}_K = \mathbf{V}_K\mathbf{V}_K^T$, where $\mathbf{V}_K \in \mathcal{O}_{d\times K}$, such that $\mathbf{V}_K^{(\ell)}\mathbf{V}_K^{(\ell)^T} = \mathbf{P}_K$ for all $\ell \in [m]$. In other words, $\{\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)^T}\}_{\ell=1}^m$ share the same top K eigenspace. Given the context of factor models, this implies that the factors have similar impact on the variation of the data across servers. Our goal now is to recover $\mathrm{Col}(\mathbf{V}_K)$ by the distributed PCA approach, namely Algorithm 1.

Recall that $\widehat{\mathbf{\Sigma}}^{(\ell)} = \frac{1}{n}\sum_{i=1}^n \mathbf{x}_i^{(\ell)}\mathbf{x}_i^{(\ell)^T}$ is the sample covariance matrix on the ℓ-th machine, and $\widehat{\mathbf{V}}_K^{(\ell)} = (\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)}) \in \mathcal{O}_{d\times K}$ stores K leading eigenvectors of $\widehat{\mathbf{\Sigma}}^{(\ell)}$. Besides, we redefine $\widetilde{\mathbf{\Sigma}} = \frac{1}{m}\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}$ and $\mathbf{\Sigma}^* = \mathbb{E}\widetilde{\mathbf{\Sigma}}$. Let $\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*, \bar{\mathbf{V}}_K^{(\ell)} \in \mathcal{O}_{d\times K}$ be the top K eigenvectors of $\widetilde{\mathbf{\Sigma}}$, $\mathbf{\Sigma}^*$ and $\mathbf{\Sigma}^{(\ell)}$, respectively. Below we present a theorem that characterizes the statistical performance of the distributed PCA under the heterogeneous factor models.

Theorem 7. *As in Theorem 6, define*

$$S^{(\ell)} := \frac{\|\mathbf{\Sigma}^{(\ell)}\|_{op}}{\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)}} \sqrt{\frac{Kr(\mathbf{\Sigma}^{(\ell)})}{N}} \quad \textit{and} \quad B^{(\ell)} := \frac{K\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2 \cdot r(\mathbf{\Sigma}^{(\ell)})}{n(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)})^2}.$$

There exist some positive constants C_1, C_2, C_3 and η such that

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq C_1 \sqrt{\frac{1}{m}\sum_{\ell=1}^{m} S^{(\ell)^2}} + \frac{C_2}{m}\sum_{\ell=1}^{m} B^{(\ell)} + C_3 \frac{\sqrt{K}}{m}\sum_{\ell=1}^{m} \frac{\|\mathbf{\Sigma}_u^{(\ell)}\|_{op}}{\lambda_K^{(\ell)}}, \tag{5.4}$$

provided that $n \geq \eta K \max_{\ell\in[m]}\{r(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2/[\lambda_K(\mathbf{\Sigma}^{(\ell)}) - \lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})]^2\}$.

The first two terms in the RHS of (5.4) are similar to those in (5.2). The third term $(\sqrt{K}/m) \cdot \sum_{\ell=1}^{m} \|\mathbf{\Sigma}_u^{(\ell)}\|_{op}/\lambda_K^{(\ell)}$ in the RHS of (5.4) characterizes the effect of heterogeneity in statistical efficiency of $\widetilde{\mathbf{V}}_K$. When $\|\mathbf{\Sigma}_u^{(\ell)}\|_{op}$ is small compared with $\lambda_K^{(\ell)}$ as in spiky factor models, $\mathbf{\Sigma}_u^{(\ell)}$ can hardly distort the eigenspace $\mathrm{Col}(\mathbf{V}_K)$ and thus has little influence on the final statistical error of $\widetilde{\mathbf{V}}_K$.

6 Simulation Study

In this section, we conduct Monte Carlo simulations to validate the statistical error rate of $\widetilde{\mathbf{V}}_K$ that is established in the previous section. We also compare the statistical accuracy of $\widetilde{\mathbf{V}}_K$ and its full sample counterpart $\widehat{\mathbf{V}}_K$, that is, the empirical top K eigenspace based on the full sample covariance. The main message is that our proposed distributed estimator performs equally well as the full sample estimator $\widehat{\mathbf{V}}_K$ when the subsample size n is large enough.

6.1 Verification of the Statistical Error Rate

Consider $\{\mathbf{x}_i\}_{i=1}^N$ i.i.d. following $N(\mathbf{0}, \mathbf{\Sigma})$, where $\mathbf{\Sigma} = \mathrm{diag}(\lambda, \lambda/2, \lambda/4, 1, \cdots, 1)$. Here the number of spiky eigenvalues $K = 3$ and $\mathbf{V}_K = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$. We generate m subsamples, each of which has n samples, and run our proposed distributed PCA algorithm (Algorithm 1) to calculate $\widetilde{\mathbf{V}}_K$. Since the centered multivariate Gaussian distribution is symmetric, according

to Theorem 4, when $\lambda = O(d)$ we have

$$\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} = O\Big(\frac{C_1\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\sqrt{\frac{Kr(\mathbf{\Sigma})}{N}}\Big) = O\Big(\sqrt{\frac{d}{mn\delta}}\Big), \tag{6.1}$$

where $\delta := \lambda_K - \lambda_{K+1} = \lambda/4 - 1$. Now we provide numerical verification of the order of the number of servers m, the eigengap δ, the subsample size n and dimension d in the statistical error.

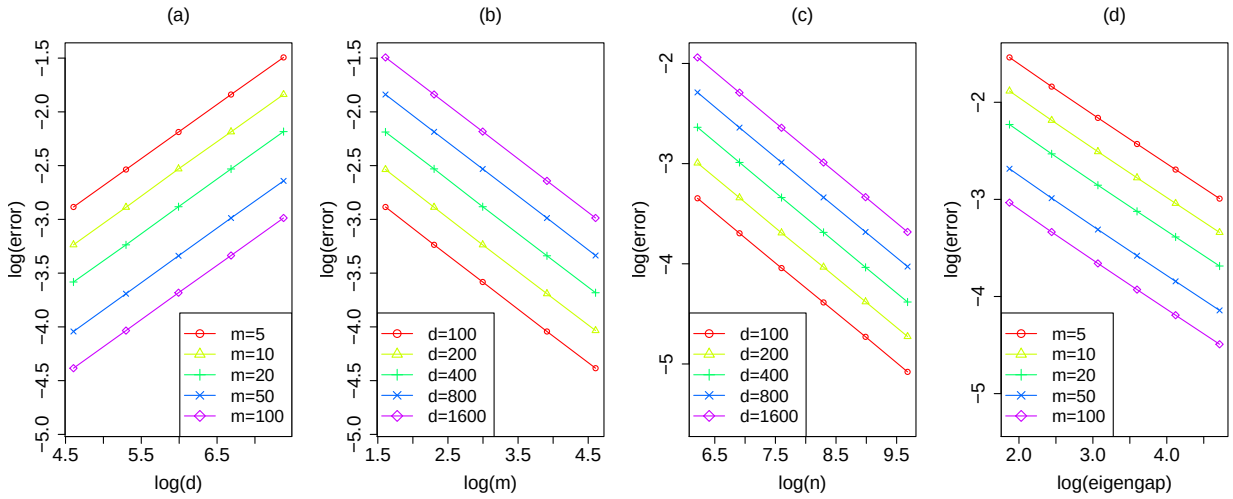

Figure 1: Statistical error rate with respect to: (a) the dimension d when $\lambda = 50$ and $n = 2000$; (b) the number of servers m when $\lambda = 50$ and $n = 2000$; (c) the subsample size n when $\lambda = 50$ and $m = 50$; (d) the eigengap δ when $d = 800$ and $n = 2000$.

Figure 1 presents four plots that demonstrate how $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ changes as d, m, n and δ increases respectively. Each data point on the plots is based on 100 independent Monte Carlo simulations. Figure 1(a) demonstrates how $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ increases with respect to the increasing dimension d when $\lambda = 50$ and $n = 2000$. Each line on the plot represents a fixed number of machines m. Figure 1(b) shows the decay rate of $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ as the number of servers m increases when $\lambda = 50$ and $n = 2000$. Different lines on the plot correspond to different dimensions d. Figure 1(c) demonstrates how $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ decays as the subsample size n increases when $\lambda = 50$ and $m = 50$. Figure 1(d) shows the relationship between $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ and the eigengap δ when $d = 800$ and $n = 2000$. The results from Figures 1(a)-1(d) show that $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ is proportion to $d^{\frac{1}{2}}$, $m^{-\frac{1}{2}}$, $n^{-\frac{1}{2}}$ and $\delta^{-\frac{1}{2}}$ respectively when the other three parameters are fixed. These empirical results are all consistent with (6.1).

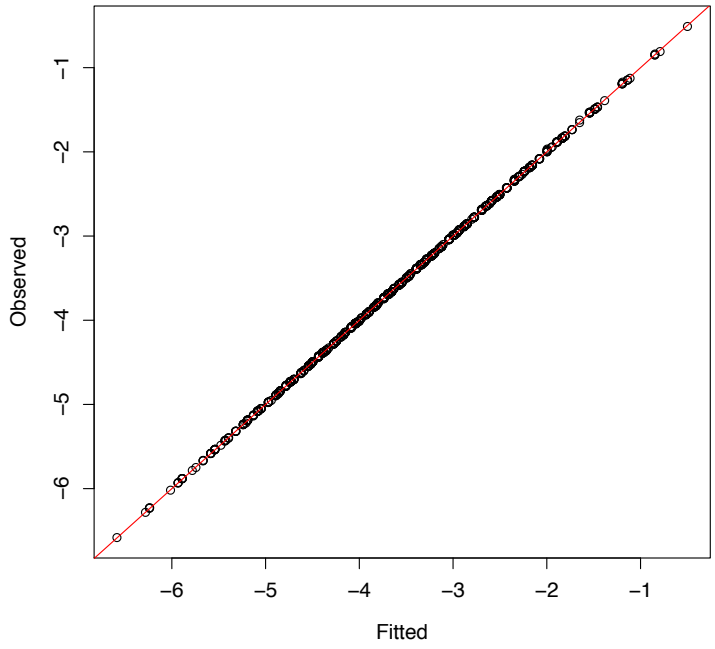

Figure 2: Observed and fitted values of $\log(\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K))$.

Figure 1 demonstrates the marginal relationship between $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ and the four parameters m, n, d and δ. Now we study their joint relationship. Inspired by (6.1), we consider a multiple regression model as follows:

$$\log(\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)) = \beta_0 + \beta_1 \log(d) + \beta_2 \log(m) + \beta_3 \log(n) + \beta_4 \log(\delta) + \varepsilon, \qquad (6.2)$$

where ε is the error term. We collect all the data points $(d, m, n, \delta, \rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K))$ from four plots in Figure 1 to fit the regression model (6.2). The fitting result is that $\widehat{\beta}_1 = 0.5043$, $\widehat{\beta}_2 = -0.4995$, $\widehat{\beta}_3 = -0.5011$ and $\widehat{\beta}_4 = -0.5120$ with the multiple $R^2 = 0.99997$. These estimates are quite consistent with the theoretical results in (6.1). Moreover, Figure 2 plots all the observed values of $\log(\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K))$ against its fitted values by the linear model (6.2). We can see that the observed and fitted values perfectly match. It indicates that the multiple regression model (6.2) well explains the joint relationship between the statistical error and the four parameters m, n, d and δ.

6.2 The Effects of Splitting

In this section we investigate how the number of data splits m affects the statistical performance of $\widetilde{\mathbf{V}}_K$ when the total sample size N is fixed. Since $N = mn$, it is easy to see that the larger m is, the smaller n will be, and hence the less computational load there will be on each individual server. In this way, to reduce the time consumption of the distributed algorithm, we prefer more splits of the data. However, per the assumptions of Theorem 4, the subsample size n should be large enough to achieve the optimal statistical performance of $\widetilde{\mathbf{V}}_K$. This motivates us to numerically illustrate how $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ changes as m increases with N fixed.

We adopt the same data generation process as described in the beginning of Section 6.1 with $\lambda = 50$ and $N = 6000$. We split the data into m subsamples where m is chosen to be all the factors of N that are less than or equal to 300. Figure 3 plots $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K)$ with respect to the number of machines m. Each point on the plot is based on 100 simulations. Each line corresponds to a different dimension d.

The results show that when the number of machines is not unreasonably large, or equivalently the number of subsample size n is not small, the statistical error does not depend on the number of machines when N is fixed. This is consistent with (6.1) where the statistical error rate only depends on the total sample size $N = mn$. When the number of machines m is large ($\log m \geq 5$), or the subsample size n is small, we observe slightly growing statistical error of the distributed PCA. This is aligned with the required lower bound of n in Theorem 4 to achieve the optimal statistical performance of $\widetilde{\mathbf{V}}_K$. Note that even when $m = 300$ ($\log(m) \approx 5.7$) and $n = 20$, our distributed PCA performs very well. This demonstrates that distributed PCA is statistically efficient as long as m is within a reasonable range.

6.3 Comparison between Distributed and Full Sample PCA

In this subsection, we compare the statistical performance of the following three methods:

1. Distributed PCA (DP)
2. Full sample PCA (FP), i.e., the PCA based on the all the samples
3. Distributed PCA with communication of five additional largest eigenvectors (DP5).

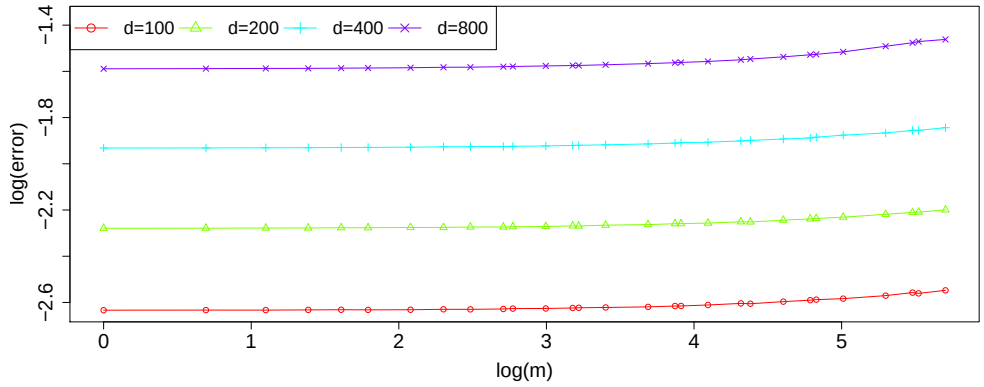

Figure 3: Statistical error with respect to the number of machines when the total sample size $N = 6000$ is fixed.

Here we explain more on the third method DP5. The difference between DP5 and DP is that on each server, DP5 calculates $\widehat{\mathbf{V}}_{K+5}^{(\ell)}$, the top $K+5$ eigenvectors of $\mathbf{\Sigma}^{(\ell)}$ and send them to the central server, and the central server computes the top K eigenvectors of $(1/m)\sum_{\ell=1}^{m}\widehat{\mathbf{V}}_{K+5}^{(\ell)}\widehat{\mathbf{V}}_{K+5}^{(\ell)^T}$ as the final output. Intuitively, DP5 communicates more information of the covariance structure and is designed to guide the spill-over effects of the eigenspace spanned by the top K eigenvalues. In Figure 4, we compare the performance of all the three methods under various scenarios.

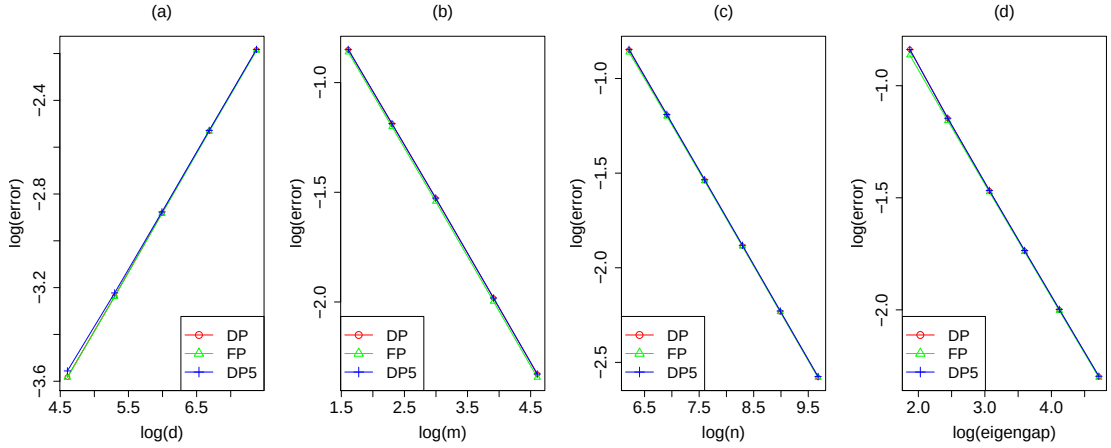

Figure 4: Comparison between DP, FP and DP5: (a) $m = 20$, $n = 2000$ and $\lambda = 50$; (b) $d = 1600$, $n = 1000$ and $\lambda = 30$; (c) $d = 800$, $m = 5$ and $\lambda = 30$; (d) $d = 1600$, $m = 10$ and $n = 500$.

From Figures 4(a)-4(d), we can see that all the three methods have similar finite sample performance. This means that it suffices to communicate K eigenvectors to enjoy the same statistical accuracy as the full sample PCA. For more challenging situations with large $p/(mn\delta)$ ratios, small improvements using FP are visible.

References

Anderson, T. W. (1963). Asymptotic theory for principal component analysis. *The Annals of Mathematical Statistics* **34** 122–148.

Baik, J., Ben Arous, G. and Péché, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability* **33** 1643–1697.

Battey, H., Fan, J., Liu, H., Lu, J. and Zhu, Z. (2015). Distributed estimation and inference with statistical guarantees. *arXiv preprint arXiv:1509.05457* .

Bertrand, A. and Moonen, M. (2014). Distributed adaptive estimation of covariance matrix eigenvectors in wireless sensor networks with application to distributed PCA. *Signal Processing* **104** 120–135.

Boutsidis, C., Woodruff, D. P. and Zhong, P. (2015). Optimal principal component analysis in distributed and streaming models. *arXiv preprint arXiv:1504.06729* .

Cai, T. T., Ma, Z. and Wu, Y. (2013). Sparse PCA: Optimal rates and adaptive estimation. *The Annals of Statistics* **41** 3074–3110.

Chen, T.-L., Chang, D. D., Huang, S.-Y., Chen, H., Lin, C. and Wang, W. (2016). Integrating multiple random sketches for singular value decomposition. *arXiv preprint arXiv:1608.08285* .

Chen, X. and Xie, M. (2014). A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica* **24** 1655–1684.

El Karoui, N. and d'Aspremont, A. (2010). Second order accurate distributed eigenvector computation for extremely large matrices. *Electronic Journal of Statistics* **4** 1345–1385.

Golub, G. H. and Van Loan, C. F. (2012). *Matrix computations*, vol. 3. JHU Press.

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology* **24** 417–441.

Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics* **29** 295–327.

Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104** 682–693.

Jung, S. and Marron, J. S. (2009). PCA consistency in high dimension, low sample size context. *The Annals of Statistics* **37** 4104–4130.

Kargupta, H., Huang, W., Sivakumar, K. and Johnson, E. (2001). Distributed clustering using collective principal component analysis. *Knowledge and Information Systems* **3** 422–448.

Kneip, A. and Utikal, K. J. (2001). Inference for density families using functional principal component analysis. *Journal of the American Statistical Association* **96** 519–542.

Koltchinskii, V. and Lounici, K. (2017). Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli* **23** 110–133.

Lee, J. D., Sun, Y., Liu, Q. and Taylor, J. E. (2015). Communication-efficient sparse regression: a one-shot approach. *arXiv preprint arXiv:1503.04337* .

Li, L., Scaglione, A. and Manton, J. H. (2011). Distributed principal subspace estimation in wireless sensor networks. *IEEE Journal of Selected Topics in Signal Processing* **5** 725–738.

Liang, Y., Balcan, M.-F. F., Kanchanapally, V. and Woodruff, D. (2014). Improved distributed principal component analysis. In *Advances in Neural Information Processing Systems*.

Nadler, B. (2008). Finite sample approximation results for principal component analysis: A matrix perturbation approach. *The Annals of Statistics* **36** 2791–2817.

Onatski, A. (2012). Asymptotics of the principal components estimator of large factor models with weakly influential factors. *Journal of Econometrics* **168** 244–258.

Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica* **17** 1617–1642.

Paul, D. and Johnstone, I. M. (2012). Augmented sparse principal component analysis for high dimensional data. *arXiv preprint arXiv:1202.1242* .

Pearson, K. (1901). On lines and planes of closest fit to systems of point in space. *Philosophical Magazine Series 6* **2** 559–572.

Qu, Y., Ostrouchov, G., Samatova, N. and Geist, A. (2002). Principal component analysis for dimension reduction in massive distributed data sets. In *IEEE International Conference on Data Mining (ICDM)*.

Reiss, M. and Wahl, M. (2016). Non-asymptotic upper bounds for the reconstruction error of PCA. *arXiv preprint arXiv:1609.03779* .

Schizas, I. D. and Aduroja, A. (2015). A distributed framework for dimensionality reduction and denoising. *IEEE Transactions on Signal Processing* **63** 6379–6394.

Shen, D., Shen, H. and Marron, J. S. (2013). Consistency of sparse PCA in high dimension, low sample size contexts. *Journal of Multivariate Analysis* **115** 317–333.

Shen, D., Shen, H., Zhu, H. and Marron, J. S. (2016). The statistics and mathematics of high dimension low sample size asymptotics. *Statistica Sinica* **26** 1747–1770.

Stewart, G. W. and Sun, J. (1990). *Matrix perturbation theory*. Academic Press.

Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027* .

Vu, V. Q. and Lei, J. (2012). Minimax rates of estimation for sparse PCA in high dimensions. In *AISTATS*, vol. 15.

Wang, W. and Fan, J. (2017). Asymptotics of empirical eigen-structure for ultra-high dimensional spiked covariance model. *The Annals of Statistics* .

Yu, Y., Wang, T. and Samworth, R. (2015). A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika* **102** 315–323.

Yurinskii, V. V. (1974). Exponential bounds for large deviations. *Theory of Probability and Its Applications* **19** 154–155.

Zhang, Y., Duchi, J. C. and Wainwright, M. J. (2013). Divide and conquer kernel ridge regression. In *COLT*.

7 Proofs and Technical Lemmas

7.1 Proof of Main Results

7.1.1 Proof of Lemma 1

Proof. To prove Lemma 1 we just need to combine the matrix concentration result in Koltchinskii and Lounici (2017), which we will present below immediately, and the Davis-Kahan Theorem.

Theorem 8. *(Koltchinskii and Lounici, 2017) Suppose* $\mathbf{x}$ *and* $\{\mathbf{x}_i\}_{i=1}^n$ *are i.i.d. sub-Gaussian random vectors in* $\mathbb{R}^d$ *with zero mean and covariance matrix* $\mathbf{\Sigma} \succeq 0$*. Let* $r(\mathbf{\Sigma}) = (\mathbb{E}\|\mathbf{x}\|_2)^2/\|\mathbf{\Sigma}\|_{op}$*. There exists a constant* $C_1 \geq 1$ *such that for* $t \geq 1$*, with probability at least* $1 - e^{-t}$*,*

$$\|\widehat{\mathbf{\Sigma}} - \mathbf{\Sigma}\|_{op} \leq C_1 \|\mathbf{\Sigma}\|_{op} \max\Big\{\sqrt{\frac{r(\mathbf{\Sigma})}{n}}, \frac{r(\mathbf{\Sigma})}{n}, \sqrt{\frac{t}{n}}, \frac{t}{n}\Big\}. \tag{7.1}$$

Combining this Theorem with a variant of Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)), we have with probability at least $1 - e^{-t}$,

$$\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K) \leq \frac{C_2\sqrt{K}\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\Big(\sqrt{\frac{r(\mathbf{\Sigma})}{n}} \vee \frac{r(\mathbf{\Sigma})}{n} \vee \sqrt{\frac{t}{n}} \vee \frac{t}{n}\Big),$$

where $C_2 > 0$ is some positive constant. According to the definition of the sub-exponential random variable (5.14) in Vershynin (2010), it is not hard to deliver from the inequality

above that

$$\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1} \leq \frac{C_2\sqrt{K}\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\sqrt{\frac{r(\mathbf{\Sigma})}{n}}. \tag{7.2}$$

Note that by Jensen's inequality,

$$\|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F \leq \mathbb{E}\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K) \leq \|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K)\|_{\psi_1}.$$

Therefore by the triangle inequality, we have with probability at least $1 - \exp(-t)$,

$$\begin{aligned}\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F &\leq \rho(\widehat{\mathbf{V}}_K^{(\ell)}, \mathbf{V}_K) + \|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F \\ &\leq \frac{2C_2\sqrt{K}\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\Big(\sqrt{\frac{r(\mathbf{\Sigma})}{n}} \vee \frac{r(\mathbf{\Sigma})}{n} \vee \sqrt{\frac{t}{n}} \vee \frac{t}{n}\Big).\end{aligned}$$

□

7.1.2 Proof of Theorem 1

Proof. By Jensen's inequality,

$$\|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F \leq \mathbb{E}\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{V}_K\mathbf{V}_K^T\|_F \leq \|\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1}. \tag{7.3}$$

When $\|\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} < 1/4$, the Weyl's inequality forces $\lambda_K(\mathbf{\Sigma}^*) > \frac{3}{4}$ and $\lambda_{K+1}(\mathbf{\Sigma}^*) < \frac{1}{4}$. Therefore by a variant of Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)),

$$\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*) \leq 4\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F. \tag{7.4}$$

In addition, by Corollary 1, there exists a constant $c' > 0$ such that

$$\begin{aligned}\Big\|(\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F)\Big\|_{\psi_1} &= \Big\|(\|\frac{1}{m}\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T} - \mathbf{\Sigma}^*\|_F)\Big\|_{\psi_1} \\ \leq \frac{c'}{\sqrt{m}}\|\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)^T} - \mathbf{\Sigma}^*\|_F\|_{\psi_1} &= \frac{c'C\|\mathbf{\Sigma}\|_{op}}{\lambda_K - \lambda_{K+1}}\sqrt{\frac{Kr(\mathbf{\Sigma})}{N}}.\end{aligned} \tag{7.5}$$

Combining this result with (7.3) achieves the conclusion of the theorem.

□

7.1.3 Proof of Theorem 2

Proof. Choose $j \in [d]$ and let $\mathbf{D} = \mathbf{I} - 2\mathbf{e}_j\mathbf{e}_j^T$. Let $\mathbf{\Sigma} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$ be the spectral decomposition of $\mathbf{\Sigma}$. Assume that $\widehat{\lambda}$ is an eigenvalue of the sample covariance $\widehat{\mathbf{\Sigma}} = (1/n)\sum_{i=1}^{n}\mathbf{x}_i\mathbf{x}_i^T$ and $\widehat{\mathbf{v}} \in \mathbb{S}^{d-1}$ is the correspondent eigenvector that satisfies $\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}} = \lambda\widehat{\mathbf{v}}$.

Define $\mathbf{z}_i = \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{V}^T\mathbf{x}_i$ and $\widehat{\mathbf{S}} = (1/n)\sum_{i=1}^{n}\mathbf{z}_i\mathbf{z}_i^T$. Note that $\widehat{\mathbf{\Sigma}} = \mathbf{V}\mathbf{\Lambda}^{\frac{1}{2}}\widehat{\mathbf{S}}\mathbf{\Lambda}^{\frac{1}{2}}\mathbf{V}^T$. Consider the matrix $\check{\mathbf{\Sigma}} = \mathbf{V}\mathbf{\Lambda}^{\frac{1}{2}}\mathbf{D}_j\widehat{\mathbf{S}}\mathbf{D}_j\mathbf{\Lambda}^{\frac{1}{2}}\mathbf{V}^T$. By the sign symmetry, $\widehat{\mathbf{\Sigma}}$ and $\check{\mathbf{\Sigma}}$ are identically distributed. It is not hard to verify that $\check{\mathbf{\Sigma}}$ also has an eigenvalue λ with the correspondent eigenvector being $\mathbf{V}\mathbf{D}\mathbf{V}^T\widehat{\mathbf{v}}$. Denote the top K eigenvectors of $\widehat{\mathbf{\Sigma}}$ by $\widehat{\mathbf{V}}_K = (\widehat{\mathbf{v}}_1, \cdots, \widehat{\mathbf{v}}_K)$ and the top K eigenvectors of $\check{\mathbf{\Sigma}}$ by $\check{\mathbf{V}}_K$. Therefore we have

$$\begin{aligned}\mathbf{V}^T\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{V} &= \mathbf{V}^T\mathbb{E}(\check{\mathbf{V}}_K\check{\mathbf{V}}_K^T)\mathbf{V} = \mathbf{V}^T\mathbf{V}\mathbf{D}_j\mathbf{V}^T\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{V}\mathbf{D}_j\mathbf{V}^T\mathbf{V} \\ &= \mathbf{D}_j\mathbf{V}^T\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{V}\mathbf{D}_j.\end{aligned}$$

Since the equation above holds for all $j \in [d]$, we can reach the conclusion that $\mathbf{V}^T\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{V}$ is diagonal, i.e, $\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)$ and $\mathbf{\Sigma}$ share the same set of eigenvectors.

When $\|\rho(\widehat{\mathbf{V}}_K, \mathbf{V}_K)\|_{\psi_1} \leq 1/4$, we have $\|\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - \mathbf{V}_K\mathbf{V}_K^T\|_F \leq 1/4$. As demonstrated above, for any $k \in [K]$, the kth column of $\mathbf{V}_K$, which we denote by $\mathbf{v}_k$, should be an eigenvector of $\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)$. Note that

$$\begin{aligned}\|\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{v}_k\|_2 &= \|(\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - \mathbf{V}_K\mathbf{V}_K^T + \mathbf{V}_K\mathbf{V}_K^T)\mathbf{v}_k\|_2 \geq -\|(\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - \mathbf{V}_K\mathbf{V}_K^T)\mathbf{v}_k\|_2 + 1 \\ &\geq 1 - \frac{1}{4} = \frac{3}{4}.\end{aligned}$$

This means that with regard to $\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)$, the correspondent eigenvalue of $\mathbf{v}_k$ is greater than $\frac{3}{4}$. Denote any eigenvector of $\mathbf{\Sigma}$ that is not in $\{\mathbf{v}_k\}_{k=1}^K$ by $\mathbf{u}$, then analogously,

$$\|\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)\mathbf{u}\|_2 = \|(\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - \mathbf{V}_K\mathbf{V}_K^T + \mathbf{V}_K\mathbf{V}_K^T)\mathbf{u}\|_2 \leq \|(\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T) - \mathbf{V}_K\mathbf{V}_K^T)\mathbf{u}\|_2 \leq \frac{1}{4}.$$

This means that for $\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)$, the correspondent eigenvalue of $\mathbf{u}$ is smaller than $\frac{1}{4}$. Therefore, the top K eigenspace of $\mathbb{E}(\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T)$ is exactly $\mathbf{V}_K$ and $\rho(\mathbf{V}_K^*, \mathbf{V}_K) = 0$. □

7.1.4 Proof of Lemma 2

Proof. The first claim follows from Theorem 2 in Yu et al. (2015) and our Lemma 5. To attack the seond one, we divide the difference:

$$\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{E}\mathbf{U}) = [\mathbf{P}_\perp \widehat{\mathbf{U}}\mathbf{H} - f(\mathbf{E}\mathbf{U})] + \mathbf{P}_\perp \widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H}) + (\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}) \tag{7.6}$$

and conquer the terms separately. Lemma 5 forces $\|\mathbf{D} - \mathbf{I}\|_F = \|\mathbf{H} - \widehat{\mathbf{H}}\|_F \leq \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T\|_F^2 \leq \varepsilon^2$. Then $\|\mathbf{P}_\perp \widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H})\|_F \leq \varepsilon^2$ and

$$\|\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F = \|\mathbf{U}\mathbf{U}^T\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F = \|\mathbf{U}\mathbf{Y}\mathbf{D}\mathbf{X}^T\mathbf{X}\mathbf{Y}^T - \mathbf{U}\|_F = \|\mathbf{U}\mathbf{Y}(\mathbf{D} - \mathbf{I})\mathbf{Y}^T\|_F \leq \varepsilon^2. \tag{7.7}$$

By triangle's inequality, it remains to show that $\|\mathbf{P}_\perp \widehat{\mathbf{U}}\mathbf{H} - f(\mathbf{E}\mathbf{U})\|_F \leq 3\varepsilon^2$. Define $\mathbf{\Lambda} = \mathrm{diag}(\lambda_{s+1}, \cdots, \lambda_{s+K})$, and $L(\mathbf{V}) = \mathbf{A}\mathbf{V} - \mathbf{V}\mathbf{\Lambda}$ for $\mathbf{V} \in \mathbb{R}^{d\times K}$. Note that $L(\mathbf{v}_1, \cdots, \mathbf{v}_K) = ((\mathbf{A} - \lambda_{s+1}\mathbf{I})\mathbf{v}_1, \cdots, (\mathbf{A} - \lambda_{s+K}\mathbf{I})\mathbf{v}_K)$, and $\mathbf{G}_j(\mathbf{A} - \lambda_{s+j}\mathbf{I}) = \mathbf{P}_\perp$ holds for all $j \in [K]$. As a result, $f(L(\mathbf{V})) = -\mathbf{P}_\perp\mathbf{V}$ for any $\mathbf{V} \in \mathbb{R}^{d\times K}$. This motivates us to work on $L(\widehat{\mathbf{U}}\mathbf{H})$ in order to study $\mathbf{P}_\perp\widehat{\mathbf{U}}\mathbf{H}$.

Let $\widehat{\mathbf{\Lambda}} = \mathrm{diag}(\widehat{\lambda}_{s+1}, \cdots, \widehat{\lambda}_{s+K})$. By definition, $\widehat{\mathbf{A}}\widehat{\mathbf{U}} = \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}}$ and

$$\begin{aligned} L(\widehat{\mathbf{U}}\mathbf{H}) &= \mathbf{A}\widehat{\mathbf{U}}\mathbf{H} - \widehat{\mathbf{U}}\mathbf{H}\mathbf{\Lambda} \\ &= (\mathbf{A} - \widehat{\mathbf{A}})\widehat{\mathbf{U}}\mathbf{H} + (\widehat{\mathbf{A}}\widehat{\mathbf{U}} - \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}})\mathbf{H} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} + \widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}) \\ &= -\mathbf{E}\widehat{\mathbf{U}}\mathbf{H} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} + \widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}). \end{aligned} \tag{7.8}$$

Now we study the images of these three terms under the linear mapping f. First, the bounded property $\|f(\cdot)\|_F \leq \Delta^{-1}\|\cdot\|_F$ implies that

$$\|f(\mathbf{E}\widehat{\mathbf{U}}\mathbf{H}) - f(\mathbf{E}\mathbf{U})\|_F \leq \Delta^{-1}\|\mathbf{E}(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\|_F \leq \Delta^{-1}\|\mathbf{E}\|_{op}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F. \tag{7.9}$$

Second, the definition of f forces $f(\mathbf{U}\mathbf{M}) = \mathbf{0}$ for all $\mathbf{M} \in \mathbb{R}^{K\times K}$. Hence the decomposition

$\widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} = \mathbf{U}\mathbf{H}^{-1}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} + (\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H}$ yields

$$\begin{aligned}
&\|f[\widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H}]\|_F = \|f[(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H}]\|_F \\
&\leq \Delta^{-1}\|(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H}\|_F \leq \Delta^{-1}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{H}^{-1}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H}\|_{op} \\
&\leq \Delta^{-1}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{H}^{-1}\|_{op}\|\mathbf{E}\|_{op}\|\mathbf{H}\|_{op} \\
&\leq \Delta^{-1}\|\mathbf{E}\|_{op}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{H}^{-1}\|_{op}.
\end{aligned} \tag{7.10}$$

Here we applied Weyl's inequality $\|\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda}\|_{op} \leq \|\mathbf{E}\|_{op}$ (Corollary IV.4.9 in Stewart and Sun (1990)) and used the fact that $\|\mathbf{H}\|_{op} = \|\widehat{\mathbf{V}}^T\mathbf{V}\|_{op} \leq 1$. Third, by similar tricks we decompose the third term

$$\widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}) = \mathbf{U}\mathbf{H}^{-1}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}) + (\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}) \tag{7.11}$$

and write

$$\begin{aligned}
&\|f[\widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda})]\|_F = \|f[(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda})]\|_F \\
&\leq \Delta^{-1}\|(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\mathbf{H}^{-1}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda})\|_F \leq \Delta^{-1}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{H}^{-1}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda})\|_{op} \\
&\leq \Delta^{-1}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{H}^{-1}\|_{op}\|\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}\|_{op}.
\end{aligned} \tag{7.12}$$

As an intermediate step, we are going to show that $\|\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}\|_{op} \leq 2\|\mathbf{E}\|_{op}$. On the one hand, $\widehat{\mathbf{A}}\widehat{\mathbf{U}} = \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}}$ yields

$$L(\widehat{\mathbf{U}}) = (\mathbf{A} - \widehat{\mathbf{A}})\widehat{\mathbf{U}} + (\widehat{\mathbf{A}}\widehat{\mathbf{U}} - \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}}) + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda}) = -\mathbf{E}\widehat{\mathbf{U}} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda}). \tag{7.13}$$

On the other hand, let $\mathbf{U}_1 = (\mathbf{u}_1, \cdots, \mathbf{u}_s, \mathbf{u}_{s+K+1}, \cdots, \mathbf{u}_d)$, $\widehat{\mathbf{U}}_1 = (\widehat{\mathbf{u}}_1, \cdots, \widehat{\mathbf{u}}_s, \widehat{\mathbf{u}}_{s+K+1}, \cdots, \widehat{\mathbf{u}}_d)$, $\mathbf{H} = \widehat{\mathbf{U}}^T\mathbf{U}$ and

$$\begin{aligned}
\mathbf{\Lambda}_1 &= \mathrm{diag}(\lambda_1, \cdots, \lambda_s, \lambda_{s+K+1}, \cdots, \lambda_d), \\
\widehat{\mathbf{\Lambda}}_1 &= \mathrm{diag}(\widehat{\lambda}_1, \cdots, \widehat{\lambda}_s, \widehat{\lambda}_{s+K+1}, \cdots, \widehat{\lambda}_d).
\end{aligned} \tag{7.14}$$

We have

$$\begin{aligned}
\mathbf{A}\widehat{\mathbf{U}} &= \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{\Lambda} & \mathbf{0}\\ \mathbf{0} & \mathbf{\Lambda}_1\end{pmatrix}\begin{pmatrix}\mathbf{U}^T\\ \mathbf{U}_1^T\end{pmatrix}\widehat{\mathbf{U}} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{\Lambda}\mathbf{H}\\ \mathbf{\Lambda}_1\mathbf{U}_1^T\widehat{\mathbf{U}}\end{pmatrix},\\
\widehat{\mathbf{U}}\mathbf{\Lambda} &= \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{U}^T\\ \mathbf{U}_1^T\end{pmatrix}\widehat{\mathbf{U}}\mathbf{\Lambda} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{H}\mathbf{\Lambda}\\ \mathbf{U}_1^T\widehat{\mathbf{U}}\mathbf{\Lambda}\end{pmatrix},
\end{aligned} \tag{7.15}$$

and as a result,

$$\|\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}\|_{op} \le \|L(\widehat{\mathbf{U}})\|_{op} = \| - \mathbf{E}\widehat{\mathbf{U}} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\|_{op} \le 2\|\mathbf{E}\|_{op}. \tag{7.16}$$

By combining (7.9), (7.10), (7.12) and (7.16), we obtain that

$$\begin{aligned}
&\|\mathbf{P}_\perp\widehat{\mathbf{U}}\mathbf{H} - f(\mathbf{E}\mathbf{U})\|_F = \| - f[L(\widehat{\mathbf{U}}\mathbf{H})] - f(\mathbf{E}\mathbf{U})\|_F\\
&\le \Delta^{-1}\|\mathbf{E}\|_{op}\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F(1 + 3\|\mathbf{H}^{-1}\|_{op}).
\end{aligned} \tag{7.17}$$

On the one hand, when $\varepsilon \le 1/\sqrt{2}$, Lemma 5 forces $\|\widehat{\mathbf{H}} - \mathbf{H}\|_F \le 1/2$ and thus $\|\mathbf{H}^{-1}\|_{op} \le 2$. On the other hand, the fact $\mathbf{U} \in \mathcal{O}_{d\times K}$ implies that $\mathbf{U}^T\mathbf{U} = \mathbf{I}_K$ and

$$\begin{aligned}
&\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F = \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{U} - \mathbf{U}\mathbf{I}_K\|_F = \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{U} - \mathbf{U}\mathbf{U}^T\mathbf{U}\|_F\\
&= \|(\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T)\mathbf{U}\|_F \le \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T\|_F \le \varepsilon.
\end{aligned} \tag{7.18}$$

Hence $\|\mathbf{P}_\perp\widehat{\mathbf{U}}\mathbf{H} - f(\mathbf{E}\mathbf{U})\|_F \le 7\varepsilon^2/\sqrt{8K} \le 3\varepsilon^2$ and we have proved the second claim in Lemma 2. Finally we come to the last one. Lemma 6 implies that

$$\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - (\mathbf{P} + \mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{P}_\perp + \mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{P})\|_F \le \sqrt{2}\|\mathbf{P}_\perp\widehat{\mathbf{U}}\|_F^2 \le \sqrt{2}\varepsilon^2. \tag{7.19}$$

Observe that

$$\begin{aligned}
&\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{P} = \mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}}[\mathbf{U}^T\mathbf{P} + (\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T\mathbf{P}]\\
&= f(\mathbf{E}\mathbf{U})\mathbf{U}^T + [\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}} - f(\mathbf{E}\mathbf{U})]\mathbf{U}^T + \mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}}(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T\mathbf{P}.
\end{aligned} \tag{7.20}$$

Also,

$$\begin{aligned}&\|[\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}}-f(\mathbf{EU})]\mathbf{U}^T\|_F\le\|\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}}-f(\mathbf{EU})\|_F\|\mathbf{U}^T\|_{op}\le 4\varepsilon^2,\\&\|\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}}(\widehat{\mathbf{U}}\widehat{\mathbf{H}}-\mathbf{U})^T\mathbf{P}\|_F=\|\mathbf{P}_\perp(\widehat{\mathbf{U}}\widehat{\mathbf{H}}-\mathbf{U})(\widehat{\mathbf{U}}\widehat{\mathbf{H}}-\mathbf{U})^T\mathbf{P}\|_F\le\|\widehat{\mathbf{U}}\widehat{\mathbf{H}}-\mathbf{U}\|_F^2\le\varepsilon^2.\end{aligned}\tag{7.21}$$

Hence $\|\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{P}-f(\mathbf{EU})\mathbf{U}^T\|_F\le 5\varepsilon^2$ and similarly, $\|\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T\mathbf{P}_\perp-\mathbf{U}f(\mathbf{EU})^T\|_F\le 5\varepsilon^2$. To sum up, we have

$$\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T-(\mathbf{U}\mathbf{U}^T+f(\mathbf{EU})\mathbf{U}^T+\mathbf{U}f(\mathbf{EU})^T)\|_F\le\sqrt{2}\varepsilon^2+10\varepsilon^2\le 12\varepsilon^2.\tag{7.22}$$

□

7.1.5 Proof of Lemma 3

Proof. Following the steps to obtain (7.3), we can derive that

$$\begin{aligned}\|\mathbf{\Sigma}^*-\mathbf{V}_K\mathbf{V}_K^T\|_F&=\|\frac{1}{m}\sum_{\ell=1}^m\mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\mathbf{V}_K\mathbf{V}_K^T\|_F\le\frac{1}{m}\sum_{\ell=1}^m\|\mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\mathbf{V}_K\mathbf{V}_K^T\|_F\\&\le\frac{1}{m}\sum_{\ell=1}^m\|\rho(\widehat{\mathbf{V}}_K^{(\ell)},\mathbf{V}_K)\|_{\psi_1}\le\frac{1}{4}.\end{aligned}$$

Therefore, the Weyl's inequality forces $\lambda_K(\mathbf{\Sigma}^*)>\frac{3}{4}$ and $\lambda_{K+1}(\mathbf{\Sigma}^*)<\frac{1}{4}$. Then by a variant of Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)),

$$\rho(\widetilde{\mathbf{V}}_K,\mathbf{V}_K^*)\le 4\|\widetilde{\mathbf{\Sigma}}-\mathbf{\Sigma}^*\|_F.\tag{7.23}$$

Given Lemma 1, an application of Theorem 6 shows that there exists a constant $C>0$ such that

$$\begin{aligned}&\Big\|(\|\widetilde{\mathbf{\Sigma}}-\mathbf{\Sigma}^*\|_F)\Big\|_{\psi_1}=\Big\|(\|\frac{1}{m}\sum_{\ell=1}^m(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\mathbb{E}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})\|_F)\Big\|_{\psi_1}\\&\le\sqrt{\frac{1}{m}\sum_{\ell=1}^m\frac{CKr(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2}{N(\lambda_K^{(\ell)}-\lambda_{K+1}^{(\ell)})^2}}.\end{aligned}$$

□

7.1.6 Proof of Theorem 3

According to Theorem 6, $\|\|\mathbf{E}\|_{op}\|_{\psi_1} \le C_1\|\mathbf{\Sigma}\|_{op}\sqrt{r(\mathbf{\Sigma})/n}$. Hence $\|\epsilon\|_{\psi_1} \le 2C_1\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}) \cdot \sqrt{2Kr(\mathbf{\Sigma})/n}$. Finally by (4.5) we have

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \le \frac{320\sqrt{2}C_1^2\|\mathbf{\Sigma}\|_{op}^2}{(\lambda_K - \lambda_{K+1})^2} \cdot \frac{Kr(\mathbf{\Sigma})}{n}.$$

7.1.7 Proof of Theorem 4

Proof. When random samples have symmetric innovations, by (7.2), there exists a constant $\eta_1 > 0$ such that as long as $n \ge \eta_1 Kr(\mathbf{\Sigma})(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}))^2$, $\rho(\mathbf{V}_K^{(\ell)}, \mathbf{V}_K) \le 1/4$. Combining Theorem 1 and Theorem 2, we can reach the conclusion.

For general distribution, again by (7.2), there exists a constant $\eta_2 > 0$ such that as long as $n \ge \eta_2 Kr(\mathbf{\Sigma})(\|\mathbf{\Sigma}\|_{op}/(\lambda_K - \lambda_{K+1}))^2$, $\rho(\mathbf{V}_K^{(\ell)}, \mathbf{V}_K) \le 1/4$. Combining Theorem 1 and Corollary 3, we can establish (4.8). Letting the second term (bias term) on the RHS of (4.8) be dominated by the first term (variance term), we can find a constant $\gamma > 0$ such that (4.9) holds if the number of servers satisfies $m \le \gamma n(\lambda_K - \lambda_{K+1})^2/(Kr(\mathbf{\Sigma})\|\mathbf{\Sigma}\|_{op}^2)$.

□

7.1.8 Proof of Theorem 5

Proof. We first focus on the first subsample $\mathbf{X}^{(1)}$ and the associated top eigenvector $\widehat{\mathbf{V}}_1^{(1)}$. For ease of notation, we temporarily drop the superscript. Let $S = \sum_{i=1}^n W_i$ and $\widehat{\mathbf{\Sigma}}_{\mathbf{Z}} = \frac{d-1}{n}\sum_{i=1}^n \mathbf{1}_{\{W_i=1\}}\mathbf{Z}_i\mathbf{Z}_i^T$. From the form of $\widehat{\mathbf{\Sigma}}$ we know that $\|\widehat{\mathbf{\Sigma}}_Z\|_{op} > (\lambda/n)\sum_{i=1}^n \mathbf{1}_{\{W_i=0\}}$ and $\|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_{op} < (\lambda/n)\sum_{i=1}^n \mathbf{1}_{\{W_i=0\}}$ lead to $\widehat{\mathbf{V}}_1 \perp \mathbf{V}_1$ and $\widehat{\mathbf{V}}_1 /\!/ \mathbf{V}_1$, respectively. Hence $P(\widehat{\mathbf{V}}_1 \perp \mathbf{V}_1) + P(\widehat{\mathbf{V}}_1 /\!/ \mathbf{V}_1) = 1$. Note that

$$\begin{aligned} \|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_N &= \mathrm{Tr}(\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}) = \frac{d-1}{n}\sum_{i=1}^n \mathbf{1}_{\{W_i=1\}} = \frac{(d-1)S}{n}, \\ \|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_{op} &\ge \frac{\|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_N}{\mathrm{rank}(\widehat{\mathbf{\Sigma}}_{\mathbf{Z}})} \ge \frac{\|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_N}{\min\{n, d-1\}} \ge \frac{(d-1)S}{n^2}. \end{aligned} \tag{7.24}$$

Then

$$
\begin{aligned}
&P(\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1) \le P\Big(\|\widehat{\mathbf{\Sigma}}_{\mathbf{Z}}\|_{op} \le \frac{\lambda}{n}(n-S)\Big) \le P\Big(\frac{(d-1)S}{n^2} \le \frac{\lambda}{n}(n-S)\Big) \\
&= P\Big(\frac{S}{n} \le \frac{1}{1+\frac{d-1}{n\lambda}}\Big) \le P\Big(\frac{S}{n} \le \frac{1}{4}\Big) = P\Big(\frac{S}{n} - \frac{1}{2} \le -\frac{1}{4}\Big) \le e^{-n/8}.
\end{aligned} \tag{7.25}
$$

Above we used the assumption $d \ge 3n\lambda + 1$ and Hoeffding's inequality. Now we finish the analysis of $\widehat{\mathbf{V}}_1^{(1)}$ and collect back the superscript.

From now on we define $S = \sum_{\ell=1}^m \mathbf{1}_{\{\widehat{\mathbf{V}}_1^{(\ell)} \not/\!/ \mathbf{V}_1\}}$. For $\widehat{\mathbf{V}}_1^{(\ell)}$, let a_ℓ be its first entry and $\mathbf{b}_\ell$ be the vector of its last $(d-1)$ entries. The dichotomy $P(\widehat{\mathbf{V}}_1^{(1)} \not/\!/ \mathbf{V}_1) + P(\widehat{\mathbf{V}}_1^{(1)} \perp \mathbf{V}_1) = 1$ mentioned above forces $|a_k| = \mathbf{1}_{\{\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1\}}$, $\|\mathbf{b}_\ell\|_2 = \mathbf{1}_{\{\widehat{\mathbf{V}}_1 \perp \mathbf{V}_1\}}$ $\widehat{\mathbf{V}}_1\widehat{\mathbf{V}}_1^T = \begin{pmatrix} \mathbf{1}_{\{\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1\}} & \mathbf{0}_{1\times(d-1)} \\ \mathbf{0}_{(d-1)\times 1} & \mathbf{b}_\ell\mathbf{b}_\ell^T \end{pmatrix}$, and

$$
\widetilde{\mathbf{\Sigma}} = \frac{1}{m}\sum_{k=1}^m \widehat{\mathbf{V}}_1\widehat{\mathbf{V}}_1^T = \begin{pmatrix} \frac{1}{m}S & \mathbf{0}_{1\times(d-1)} \\ \mathbf{0}_{(d-1)\times 1} & \frac{1}{m}\sum_{k=1}^m \mathbf{b}_\ell\mathbf{b}_\ell^T \end{pmatrix}. \tag{7.26}
$$

By the assumptions $P(\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1) \le e^{-n/8} \le d^{-4}$.

Case 1: $m \le d^3$

In this case, $P(S=0) = [1 - P(\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1)]^m \ge 1 - mP(\widehat{\mathbf{V}}_1 \not/\!/ \mathbf{V}_1) \ge 1 - d^{-1}$. When $S = 0$, we have $\|\mathbf{b}_\ell\|_2 = 1$ for all $\ell \in [m]$ and $\|(1/m)\sum_{\ell=1}^m \mathbf{b}_\ell\mathbf{b}_\ell^T\|_{op} > 0$, leading to $\widetilde{\mathbf{V}}_1 \perp \mathbf{V}_1$.

Case 2: $m > d^3$

On the one hand, by Hoeffding's inequality we obtain

$$
P\Big(\frac{S}{m} \ge \frac{1}{d}\Big) \le P\Big(\frac{1}{m}(S - \mathbb{E}S) \ge \frac{1}{2d}\Big) \le e^{-2m(\frac{1}{2d})^2} = e^{-\frac{m}{2d^2}} < e^{-d/2}. \tag{7.27}
$$

On the other hand, note that

$$
\begin{aligned}
&\Big\|\frac{1}{m}\sum_{k=1}^m \mathbf{b}_\ell\mathbf{b}_\ell^T\Big\|_N = \frac{1}{m}\sum_{k=1}^m \|\mathbf{b}_\ell\|_2^2 = 1 - \frac{S}{m}, \\
&\Big\|\frac{1}{m}\sum_{k=1}^m \mathbf{b}_\ell\mathbf{b}_\ell^T\Big\|_{op} \ge \frac{1}{d-1}\Big\|\frac{1}{m}\sum_{k=1}^m \mathbf{b}_\ell\mathbf{b}_\ell^T\Big\|_N = \frac{1}{d-1}\Big(1 - \frac{S}{m}\Big).
\end{aligned} \tag{7.28}
$$

Hence

$$P(\widetilde{\mathbf{V}}_1 \perp \mathbf{V}_1) \geq P\Big(\Big\|\frac{1}{m}\sum_{k=1}^{m}\mathbf{b}_\ell\mathbf{b}_\ell^T\Big\|_{op} > \frac{S}{m}\Big) \geq P\Big[\frac{1}{d-1}\Big(1-\frac{S}{m}\Big) > \frac{S}{m}\Big] = P\Big(\frac{S}{m} < \frac{1}{d}\Big) \geq 1-e^{-d/2}. \tag{7.29}$$

□

7.1.9 Proof of Theorem 7

Proof. The proof is divided into two parts concerning the variance and bias respectively.

1. (Variance) There exists constant $C_1 > 0$ such that $\|\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*)\|_{\psi_1} \leq C_1\sqrt{\frac{1}{m}\sum_{\ell=1}^{m} S^{(\ell)^2}}$;

2. (Bias) There exist constants $C_2, C_3 > 0$ such that

$$\rho(\mathbf{V}_K^*, \mathbf{V}_K) \leq \frac{C_2}{m}\sum_{\ell=1}^{m} B^{(\ell)} + C_3\frac{\sqrt{K}}{m}\sum_{\ell=1}^{m}\frac{\|\mathbf{\Sigma}_u^{(\ell)}\|_{op}}{\lambda_K^{(\ell)}}.$$

We work on the variance part first. Let $\mathbf{\Sigma}^{(\ell)} = \bar{\mathbf{V}}^{(\ell)}\bar{\mathbf{\Lambda}}^{(\ell)}\bar{\mathbf{V}}^{(\ell)^T}$ be the spectral decomposition of $\mathbf{\Sigma}^{(\ell)}$. Note that $\bar{\mathbf{V}}_K^{(\ell)}$ is different from $\mathbf{V}_K^{(\ell)}$. Then it follows that $\Big\|\|\widehat{\mathbf{\Sigma}}^{(\ell)} - \mathbf{\Sigma}^{(\ell)}\|_{op}\Big\|_{\psi_1} \leq C\|\mathbf{\Sigma}^{(\ell)}\|_{op}\sqrt{r(\mathbf{\Sigma}^{(\ell)})/n}$ for some $C > 0$, and

$$\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \bar{\mathbf{V}}_K^{(\ell)}) \leq \sqrt{8K}\|\widehat{\mathbf{\Sigma}}^{(\ell)} - \mathbf{\Sigma}^{(\ell)}\|_{op}/(\lambda_K(\mathbf{\Sigma}^{(\ell)}) - \lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})).$$

There exists $\eta > 0$ such that $n \geq \eta K \max_{\ell\in[m]}\{r(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2/[\lambda_K(\mathbf{\Sigma}^{(\ell)}) - \lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})]^2\}$ implies $\max_{\ell\in[m]}\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \bar{\mathbf{V}}_K^{(\ell)})\|_{\psi_1} \leq 1/4$ and thus,

$$\begin{aligned}\|\mathbf{\Sigma}^* - \frac{1}{m}\sum_{\ell=1}^{m}\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\|_F &= \|\frac{1}{m}\sum_{\ell=1}^{m}\big(\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}) - \bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\big)\|_F \\ &\leq \frac{1}{m}\sum_{\ell=1}^{m}\|\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}) - \bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\|_F \leq \frac{1}{m}\sum_{\ell=1}^{m}\|\rho(\widehat{\mathbf{V}}_K^{(\ell)}, \bar{\mathbf{V}}_K^{(\ell)^T})\|_{\psi_1} \leq \frac{1}{4}.\end{aligned}$$

Therefore, the Weyl's inequality forces $\lambda_K(\mathbf{\Sigma}^*) > \frac{3}{4}$ and $\lambda_{K+1}(\mathbf{\Sigma}^*) < \frac{1}{4}$. Then by a variant of Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)), $\rho(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*) \leq 4\|\widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\|_F$. Given Lemma 1, an application of Corollary 1 shows that there exist constants $C_1, C_1', C_1'' > 0$ such

that

$$\begin{aligned}
&\Big\|(\|\widetilde{\mathbf{\Sigma}}-\mathbf{\Sigma}^*\|_F)\Big\|_{\psi_1} = \Big\|\|\frac{1}{m}\sum_{\ell=1}^m[\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})]\|_F\Big\|_{\psi_1} \\
&\le \sqrt{\frac{C_1'}{m}\sum_{\ell=1}^m\Big\|\|[\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})]\|_F\Big\|_{\psi_1}^2} \\
&\le \frac{C_1''}{m}\sqrt{\sum_{\ell=1}^m\Big\|\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T}-\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\|_F\Big\|_{\psi_1}^2} \le \frac{C_1''}{m}\sqrt{\sum_{\ell=1}^m\Big\|\frac{\sqrt{K}\|\widehat{\mathbf{\Sigma}}^{(\ell)}-\mathbf{\Sigma}^{(\ell)}\|_{op}}{\lambda_K(\mathbf{\Sigma}^{(\ell)})-\lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})}\Big\|_{\psi_1}^2} \\
&\le \frac{C_1}{4m}\sqrt{\sum_{\ell=1}^m\frac{Kr(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2}{n[\lambda_K(\mathbf{\Sigma}^{(\ell)})-\lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})]^2}} = \frac{C_1}{4}\sqrt{\frac{1}{m}\sum_{\ell=1}^m S^{(\ell)^2}}.
\end{aligned}$$

Now we come to the bias part. Again as we derive in (4.5), we have

$$\begin{aligned}
&\rho(\mathbf{V}_K^*,\mathbf{V}_K)\le 2\sqrt{2}\cdot\|\mathbf{\Sigma}^*-\mathbf{V}_K\mathbf{V}_K^T\|_F \le \frac{2\sqrt{2}}{m}\sum_{\ell=1}^m\|\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})-\mathbf{V}_K\mathbf{V}_K^T\|_F \\
&\le \frac{2\sqrt{2}}{m}\sum_{\ell=1}^m\|\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})-\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\|_F + \frac{2\sqrt{2}}{m}\sum_{\ell=1}^m\|\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}-\mathbf{V}_K\mathbf{V}_K^T\|_F
\end{aligned} \tag{7.30}$$

We apply Lemma 2 to $\widehat{\mathbf{\Sigma}}^{(\ell)}$ and $\mathbf{\Sigma}^{(\ell)}$. There exists $C_2', C_2'' > 0$ such that for $\ell \in [m]$,

$$\|\mathbb{E}(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)^T})-\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}\|_F \le C_2'\Big(\frac{\sqrt{K}\|(\|\widehat{\mathbf{\Sigma}}^{(\ell)}-\mathbf{\Sigma}^{(\ell)}\|_{op})\|_{\psi_1}}{\lambda_K(\mathbf{\Sigma}^{(\ell)})-\lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})}\Big)^2 \le \frac{C_2''Kr(\mathbf{\Sigma}^{(\ell)})\|\mathbf{\Sigma}^{(\ell)}\|_{op}^2}{n[\lambda_K(\mathbf{\Sigma}^{(\ell)})-\lambda_{K+1}(\mathbf{\Sigma}^{(\ell)})]^2}.$$

Applying Theorem 2 in Yu et al. (2015) to $\mathbf{\Sigma}^{(\ell)}$ and $\mathbf{V}_K^{(\ell)}\mathbf{\Lambda}_K^{(\ell)}\mathbf{V}_K^{(\ell)^T}$, we obtain that

$$\|\bar{\mathbf{V}}_K^{(\ell)}\bar{\mathbf{V}}_K^{(\ell)^T}-\mathbf{V}_K\mathbf{V}_K^T\|_F \le \frac{2\sqrt{K}\|\mathbf{\Sigma}_u^{(\ell)}\|_{op}}{\lambda_K^{(\ell)}}.$$

Then the proof is completed.

□

7.2 Technical lemmas

7.2.1 Tail bounds

The following result is a Bernstein-type inequality for random vectors in Banach spaces whose norms are sub-exponential, a variant of the main theorem in Yurinskii's note Yurinskii (1974). In Yurinskii's original result, the second order moment was used to control the moment growth, similar to the definition of sub-gaussian vectors in Koltchinskii and Lounici (2017). Our result is only concerned about the sub-exponential norms and does not involve second order moments. This relaxation is crucial when aggregating PCA estimators, since it is natural to bound errors of individual estimators from above, but unnatural and hard to bound squared errors from below.

Lemma 4. *If the independent random vector $\mathbf{u}_1, \cdots, \mathbf{u}_n$ with values in a separable Banach space (where the norm is denoted by $\|\cdot\|$) satisfy the conditions*

$$\mathbb{E}\mathbf{u}_i = \mathbf{0}, \quad \mathbb{E}\|\mathbf{u}_i\|^k \leq \frac{k!}{2} L^k, \quad \forall k \geq 2, \quad i \in [n], \tag{7.31}$$

then

$$P(\|n^{-1}\mathbf{w}\| \geq t) \leq \exp\Big(- \frac{nt^2 - 2\beta t}{4L(t+2L)} \Big), \tag{7.32}$$

where $\mathbf{w} = \sum_{i=1}^n \mathbf{u}_i$ and $\beta = \mathbb{E}\|\mathbf{w}\|$.

Proof. The proof has the same essence as the one in Yurinskii (1974). The only difference lies in the moment assumption that is used at the end of the proof. For $t \geq 0$ and $h \geq 0$, we have

$$P(\|n^{-1}\mathbf{w}\| \geq t) = P(\|\mathbf{w}\| \geq nt) \leq e^{-nht} \mathbb{E} e^{h\|\mathbf{w}\|}. \tag{7.33}$$

We will prove the concentration inequality by estimating $\mathbb{E}e^{h\|\mathbf{w}\|}$. Let $(\Omega, \mathcal{F}, P)$ denote the probability space where $\{\mathbf{u}_i\}_{i=1}^n$ are defined. For $i \in [n]$ we define $\mathbf{v}_i = \mathbf{w} - \mathbf{u}_i$ and $\mathcal{F}_i = \sigma(\mathbf{u}_1, \cdots, \mathbf{u}_i)$. Then let $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and

$$\begin{aligned} &\mathbb{E}_i(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_{i-1}), \quad i \in [n+1], \\ &\zeta_i = \mathbb{E}_{i+1}\|\mathbf{w}\| - \mathbb{E}_i\|\mathbf{w}\|, \quad i \in [n]. \end{aligned} \tag{7.34}$$

Suppose we have $\mathbb{E}_i e^{h\zeta_i} \le f_i(h)$ for $i \in [n]$ and $h \ge 0$. For $2 \le i \le n+1$, we observe that

$$\begin{aligned}\mathbb{E}e^{h\mathbb{E}_i\|\mathbf{w}\|} &= \mathbb{E}(\mathbb{E}_{i-1}e^{h\mathbb{E}_i\|\mathbf{w}\|}) = \mathbb{E}(\mathbb{E}_{i-1}e^{h\zeta_{i-1}+h\mathbb{E}_{i-1}\|\mathbf{w}\|}) \\ &= \mathbb{E}(e^{h\mathbb{E}_{i-1}\|\mathbf{w}\|}\cdot\mathbb{E}_{i-1}e^{h\zeta_{i-1}}) \le f_{i-1}(h)\mathbb{E}e^{h\mathbb{E}_{i-1}\|\mathbf{w}\|}.\end{aligned} \tag{7.35}$$

By induction, it follows that $\mathbb{E}e^{h\|\mathbf{w}\|} \le e^{h\beta}\prod_{i=1}^n f_i(h)$. Then we come to find upper bounds $f_i(h)$ for $\mathbb{E}_i e^{h\zeta_i}$. On the one hand,

$$\begin{aligned}\zeta_i &= \mathbb{E}_{i+1}\|\mathbf{w}\| - \mathbb{E}_i\|\mathbf{w}\| = \mathbb{E}_{i+1}\|\mathbf{u}_i+\mathbf{v}_i\| - \mathbb{E}_i\|\mathbf{u}_i+\mathbf{v}_i\| \\ &\ge (\mathbb{E}_{i+1}\|\mathbf{v}_i\| - \mathbb{E}_{i+1}\|\mathbf{u}_i\|) - (\mathbb{E}_i\|\mathbf{v}_i\| + \mathbb{E}_i\|\mathbf{u}_i\|) = -(\|\mathbf{u}_i\| + \mathbb{E}\|\mathbf{u}_i\|),\end{aligned} \tag{7.36}$$

where the last equality is due to $\mathbb{E}_{i+1}\|\mathbf{v}_i\| = \mathbb{E}(W-\mathbf{u}_i|\mathcal{F}_i) = \mathbb{E}(W-\mathbf{u}_i|\mathcal{F}_{i-1}) = \mathbb{E}_i\|\mathbf{v}_i\|$. On the other hand,

$$\begin{aligned}&\zeta_i = \mathbb{E}_{i+1}\|\mathbf{w}\| - \mathbb{E}_i\|\mathbf{w}\| = \mathbb{E}_{i+1}\|\mathbf{u}_i+\mathbf{v}_i\| - \mathbb{E}_i\|\mathbf{u}_i+\mathbf{v}_i\| \\ &\le (\mathbb{E}_{i+1}\|\mathbf{v}_i\| + \mathbb{E}_{i+1}\|\mathbf{u}_i\|) - (\mathbb{E}_i\|\mathbf{v}_i\| - \mathbb{E}_i\|\mathbf{u}_i\|) = \mathbb{E}_{i+1}\|\mathbf{u}_i\| + \mathbb{E}_i\|\mathbf{u}_i\| = \|\mathbf{u}_i\| + \mathbb{E}\|\mathbf{u}_i\|.\end{aligned} \tag{7.37}$$

Thus $|\zeta_i| \le \|\mathbf{u}_i\| + \mathbb{E}\|\mathbf{u}_i\|$. When $k \ge 1$,

$$\begin{aligned}|\zeta_i|^k &\le (\|\mathbf{u}_i\| + \mathbb{E}\|\mathbf{u}_i\|)^k = [2(\frac{1}{2}\|\mathbf{u}_i\| + \frac{1}{2}\mathbb{E}\|\mathbf{u}_i\|)]^k \\ &\le 2^k(\frac{1}{2}\|\mathbf{u}_i\|^k + \frac{1}{2}\mathbb{E}\|\mathbf{u}_i\|^k) = 2^{k-1}(\|\mathbf{u}_i\|^k + \mathbb{E}\|\mathbf{u}_i\|^k)\end{aligned} \tag{7.38}$$

holds by Jensen's inequality. Together with the moment assumption, we deduce that

$$\mathbb{E}_i|\zeta_i|^k \le 2^{k-1}(\mathbb{E}_i\|\mathbf{u}_i\|^k + \mathbb{E}_i\mathbb{E}\|\mathbf{u}_i\|^k) = 2^k\mathbb{E}\|\mathbf{u}_i\|^k \le \frac{k!}{2}(2L)^k. \tag{7.39}$$

Note that $\mathbb{E}_i\zeta_i = 0$ almost surely holds. Then

$$\begin{aligned}\mathbb{E}_i e^{h\zeta_i} &= \sum_{i=0}^\infty \frac{h^k}{k!}\mathbb{E}_i\zeta_i^k \le 1 + \sum_{i=2}^\infty \frac{h^k}{k!}\mathbb{E}_i|\zeta_i|^k \le 1 + \sum_{i=2}^\infty \frac{h^k}{k!}\cdot\frac{k!}{2}(2L)^k \\ &= 1 + \frac{2L^2h^2}{1-2Lh} \le \exp\Big(\frac{2L^2h^2}{1-2Lh}\Big) \triangleq f_i(h), \quad 0 \le h < \frac{1}{2L}.\end{aligned} \tag{7.40}$$

As a result,

$$\mathbb{E}e^{h\|\mathbf{w}\|} \leq \exp\Big(h\beta + \frac{2nL^2h^2}{1-2Lh}\Big), \quad 0 \leq h < \frac{1}{2L}. \tag{7.41}$$

The final result follows by taking $h = (2L + 4L^2/t)^{-1}$. □

We will use a special case of Lemma 4, where the random vectors take values in a separable Hilbert space.

Corollary 1. *If random vectors $\mathbf{u}_1, \cdots, \mathbf{u}_n$ with values in a separable Hilbert space (where the norm is denoted by $\|\cdot\|$) satisfy the conditions*

$$\mathbb{E}\mathbf{u}_i = \mathbf{0}, \quad \|\|\mathbf{u}_i\|\|_{\psi_1} \leq L_i, \quad \forall k \geq 2, \quad i \in [n], \tag{7.42}$$

then

$$P(\|\frac{1}{n}\sum_{i=1}^{n}\mathbf{u}_i\| \geq t) \leq \exp\Big(1 - \frac{n}{8}\min\Big\{\frac{t}{2L}, \Big(\frac{t}{2L}\Big)^2\Big\}\Big), \quad \forall t \geq 0, \quad n \geq 4, \tag{7.43}$$

where $L = \sqrt{\sum_{i=1}^{n} L_i^2/n}$. Furthermore, there exists a universal constant $c' > 0$ such that $\|\|(1/n)\sum_{i=1}^{n}\mathbf{u}_i\|\|_{\psi_1} \leq c'L/\sqrt{n}$.

Proof. By independence, orthogonality and the moment assumption,

$$\beta^2 = \mathbb{E}^2\|\sum_{i=1}^{n}\mathbf{u}_i\| \leq \mathbb{E}\|\sum_{i=1}^{n}\mathbf{u}_i\|^2 = \sum_{i=1}^{n}\mathbb{E}\|\mathbf{u}_i\|^2 \leq \sum_{i=1}^{n}\frac{2!}{2}L_i^2 = nL^2, \tag{7.44}$$

where we write $\sqrt{\sum_{i=1}^{n} L_i^2/n}$ as L. Hence $\beta \leq \sqrt{n}L$. By plugging this into Lemma 4, we derive that

$$P(\|\frac{1}{n}\sum_{i=1}^{n}\mathbf{u}_i\| \geq t) \leq \exp\Big(-\frac{nt^2 - 2\sqrt{n}Lt}{4L(t+2L)}\Big), \quad \forall t \geq 0. \tag{7.45}$$

From now on we assume $n \geq 4$. When $t > 2L$, we have $4L(t+2L) < 4L(t+t) = 8Lt$ and

$$nt^2 - 2\sqrt{n}Lt > nt^2 - 2\sqrt{n}\cdot\frac{t}{2}\cdot t = \sqrt{n}t^2(\sqrt{n}-1) \geq \frac{n}{2}t^2, \quad \forall n \geq 4. \tag{7.46}$$

So

$$\frac{nt^2 - 2\sqrt{n}Lt}{4L(t+2L)} \geq \frac{nt}{16L}, \quad \forall n \geq 4,\ t > 2L. \tag{7.47}$$

When $\frac{4L}{\sqrt{n}} \leq t \leq 2L$, we have $4L(t+2L) \leq 4L(2L+2L) = 16L^2$ and

$$nt^2 - 2\sqrt{n}Lt = \sqrt{n}t(\sqrt{n}t - 2L) \geq \sqrt{n}t(\sqrt{n}t - \frac{\sqrt{n}t}{2}) \geq \frac{nt^2}{2}, \quad \forall n \geq 4. \tag{7.48}$$

So

$$\frac{nt^2 - 2\sqrt{n}Lt}{4L(t+2L)} \geq \frac{nt^2}{32L^2}, \quad \forall n \geq 4,\ \frac{4L}{\sqrt{n}} \leq t \leq 2L. \tag{7.49}$$

When $0 \leq t < \frac{4L}{\sqrt{n}}$, we have the trivial estimate

$$P(\|\frac{1}{n}\sum_{i=1}^{n}\mathbf{u}_i\| \geq t) \leq 1 < \exp\Big(1 - \frac{nt^2}{32L^2}\Big). \tag{7.50}$$

The desired result follows from all the estimates above.

Now it remains to show that $\|\|(1/n)\sum_{i=1}^{n}\mathbf{u}_i\|\|_{\psi_1} \leq c'L/\sqrt{n}$. First we claim that

$$P(\|(1/n)\sum_{i=1}^{n}\mathbf{u}_i\| \geq t) \leq \exp\big(1 - \frac{\sqrt{n}t}{8L}\big) \tag{7.51}$$

for all $t > 0$. Note that when $t \leq 8L/\sqrt{n}$, (7.51) holds automatically because the RHS is larger than 1. Therefore we only need to verify (7.51) when $t > 8L/\sqrt{n}$. It is not hard to verify that when $t > 4L/\sqrt{n}$ and $n \geq 4$,

$$\exp\Big(1 - \frac{n}{8}\min\Big\{\frac{t}{2L}, \Big(\frac{t}{2L}\Big)^2\Big\}\Big) \leq \exp\big(1 - \frac{\sqrt{n}t}{8L}\big).$$

Then the proof is completed by the equivalent definitions of sub-exponential norm, see Definition 5.13 of Vershynin (2010). □

7.2.2 Matrix analysis

Let $\mathbf{V}, \widehat{\mathbf{V}} \in \mathcal{O}_{d\times K}$, and use $\{\sigma_j\}_{j=1}^{K}$ to denote the singular values of $\mathbf{H} = \widehat{\mathbf{V}}^T\mathbf{V}$. There are K principal angles $\{\theta_j\}_{j=1}^{K}$ between $\mathrm{Col}(\mathbf{V})$ and $\mathrm{Col}(\widehat{\mathbf{V}})$, defined as $\theta_j = \cos^{-1}\sigma_j$. Then we let $\sin\boldsymbol{\Theta}(\widehat{\mathbf{V}}, \mathbf{V}) = \mathrm{diag}(\sin\theta_1, \cdots, \sin\theta_K)$. The following lemma presents the relationship

among different subspace distances. It also investigates the optimal orthonormal transform $\widehat{\mathbf{H}}$ minimizing $\|\widehat{\mathbf{V}}\widehat{\mathbf{H}} - \mathbf{V}\|_F$, i.e. the solution to the Procrustes problem.

Lemma 5. *Let* $\mathbf{V}, \widehat{\mathbf{V}} \in \mathcal{O}_{d\times K}$, $\mathbf{H} = \widehat{\mathbf{V}}^T\mathbf{V}$, $\mathbf{H} = \mathbf{X}\mathbf{D}\mathbf{Y}^T$ *be its SVD, and* $\widehat{\mathbf{H}} = \mathbf{X}\mathbf{Y}^T$. *We have*

$$\begin{aligned}&\|\widehat{\mathbf{H}} - \mathbf{H}\|_F^{1/2} \le \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\widehat{\mathbf{V}}\|_F \le \min_{\mathbf{O}\in\mathcal{O}_{K,K}} \|\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V}\|_F \\ &= \|\widehat{\mathbf{V}}\widehat{\mathbf{H}} - \mathbf{V}\|_F \le \sqrt{2}\|\sin\boldsymbol{\Theta}(\widehat{\mathbf{V}}, \mathbf{V})\|_F = \|\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T - \mathbf{V}\mathbf{V}^T\|_F.\end{aligned} \tag{7.52}$$

Proof. Obviously, $0 \le \mathbf{D}_{ii} \le 1$ for all $i \in [K]$, and $\|\widehat{\mathbf{H}} - \mathbf{H}\|_F = \|\mathbf{X}(\mathbf{D} - \mathbf{I})\mathbf{Y}^T\|_F = \|\mathbf{D} - \mathbf{I}\|_F$. Let $\mathbf{P} = \mathbf{V}\mathbf{V}^T$. By Lemma 7, we obtain that

$$\begin{aligned}&\|(\mathbf{I} - \mathbf{P})\widehat{\mathbf{V}}\|_F^2 \ge \|\mathbf{P}\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T\mathbf{P} - \mathbf{P}\|_F = \|\mathbf{V}\mathbf{H}^T\mathbf{H}\mathbf{V}^T - \mathbf{V}\mathbf{V}^T\|_F = \|\mathbf{V}\mathbf{Y}\mathbf{D}^2\mathbf{Y}^T\mathbf{V}^T - \mathbf{V}\mathbf{V}^T\|_F \\ &= \|(\mathbf{V}\mathbf{Y})(\mathbf{D}^2 - \mathbf{I})(\mathbf{V}\mathbf{Y})^T\|_F = \|\mathbf{D}^2 - \mathbf{I}\|_F \ge \|\mathbf{D} - \mathbf{I}\|_F.\end{aligned} \tag{7.53}$$

Hence $\|\widehat{\mathbf{H}} - \mathbf{H}\|_F^{1/2} \le \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\widehat{\mathbf{V}}\|_F$. For any $\mathbf{O} \in \mathcal{O}_{K,K}$, we have

$$\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\widehat{\mathbf{V}}\|_F = \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\widehat{\mathbf{V}}\mathbf{O}\|_F = \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)(\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V})\|_F \le \|\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V}\|_F. \tag{7.54}$$

$\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\widehat{\mathbf{V}}\|_F \le \min_{\mathbf{O}\in\mathcal{O}_{K,K}} \|\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V}\|_F$ is derived by taking infimum. Again, for any $\mathbf{O} \in \mathcal{O}_{K,K}$ we have

$$\begin{aligned}&\|\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V}\|_F^2 = \|\widehat{\mathbf{V}}\mathbf{O}\|_F^2 + \|\mathbf{V}\|_F^2 - 2\mathrm{Tr}(\mathbf{O}^T\widehat{\mathbf{V}}^T\mathbf{V}) = 2K - 2\mathrm{Tr}(\mathbf{O}^T\mathbf{H}) \\ &\ge 2K - 2\|\mathbf{O}\|_{op}\|\mathbf{H}\|_* = 2K - 2\mathrm{Tr}(\widehat{\mathbf{H}}^T\mathbf{H}) = \|\widehat{\mathbf{V}}\widehat{\mathbf{H}} - \mathbf{V}\|_F^2.\end{aligned} \tag{7.55}$$

This proves $\min_{\mathbf{O}\in\mathcal{O}_{K,K}} \|\widehat{\mathbf{V}}\mathbf{O} - \mathbf{V}\|_F = \|\widehat{\mathbf{V}}\widehat{\mathbf{H}} - \mathbf{V}\|_F$. From above we also see that

$$\|\widehat{\mathbf{V}}\widehat{\mathbf{H}} - \mathbf{V}\|_F^2 = 2\sum_{j=1}^K (1 - \mathbf{D}_{jj}) \le 2\sum_{j=1}^K (1 - \mathbf{D}_{jj}^2) = 2\|\sin\boldsymbol{\Theta}(\widehat{\mathbf{V}}, \mathbf{V})\|_F^2. \tag{7.56}$$

Besides,

$$\begin{aligned}\|\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T - \mathbf{V}\mathbf{V}^T\|_F^2 &= \|\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T\|_F^2 + \|\mathbf{V}\mathbf{V}^T\|_F^2 - 2\mathrm{Tr}(\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T\mathbf{V}\mathbf{V}^T) \\ &= 2K - 2\|\mathbf{H}\|_F^2 = 2\sum_{j=1}^{K}(1 - \mathbf{D}_{jj}^2) = 2\|\sin\boldsymbol{\Theta}(\widehat{\mathbf{V}}, \mathbf{V})\|_F^2.\end{aligned} \tag{7.57}$$

Then the proof is completed by combining the estimates above. □

Lemma 6. *Suppose $\mathbf{V} \in \mathcal{O}_{d\times K}$ and $\mathbf{P} \in \mathbb{R}^{d\times d}$ is a projection matrix with rank K. Then*

$$\|(\mathbf{V}\mathbf{V}^T - \mathbf{P}) - (\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp + \mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P})\|_F \le \sqrt{2}\|(\mathbf{I}_d - \mathbf{P})\mathbf{V}\|_F^2. \tag{7.58}$$

Proof. Let $\mathbf{P}_\perp = \mathbf{I}_d - \mathbf{P}$. We have

$$\begin{aligned}\mathbf{V}\mathbf{V}^T - \mathbf{P} &= (\mathbf{P} + \mathbf{P}_\perp)\mathbf{V}\mathbf{V}^T(\mathbf{P} + \mathbf{P}_\perp) - \mathbf{P} \\ &= (\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P} - \mathbf{P}) + (\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp + \mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P}) + \mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp.\end{aligned} \tag{7.59}$$

Therefore,

$$\begin{aligned}&\|(\mathbf{V}\mathbf{V}^T - \mathbf{P}) - (\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp + \mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P})\|_F^2 = \|(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P} - \mathbf{P}) + \mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp\|_F^2 \\ &= \|\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P} - \mathbf{P}\|_F^2 + \|\mathbf{P}_\perp\mathbf{V}\mathbf{V}^T\mathbf{P}_\perp\|_F^2 \\ &\le 2\|(\mathbf{I}_d - \mathbf{P})\mathbf{V}\|_F^4.\end{aligned} \tag{7.60}$$

The inequality follows from Lemma 7 and submultiplicativity of $\|\cdot\|_F$. □

Lemma 7. *Suppose $V \in \mathcal{O}_{d\times K}$ and $P \in \mathbb{R}^{d\times d}$ is a projection matrix with rank K. Then* $\|\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P} - \mathbf{P}\|_F \le \|(\mathbf{I}_d - \mathbf{P})\mathbf{V}\|_F^2$.

Proof. Since $\mathbf{V}\mathbf{V}^T$ is a projection matrix with rank K, Lemma 8 tells us that

$$\begin{aligned}&\|\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P} - \mathbf{P}\|_F^2 = \sum_{j=1}^{K}[1 - \sigma_j(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P})]^2 \le \Big(\sum_{j=1}^{K}|1 - \sigma_j(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P})|\Big)^2 \\ &= \Big(\sum_{j=1}^{K}[1 - \sigma_j(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P})]\Big)^2 = \Big(K - \sum_{j=1}^{K}\sigma_j(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P})\Big)^2 \\ &= [K - \mathrm{Tr}(\mathbf{P}\mathbf{V}\mathbf{V}^T\mathbf{P})]^2 = (K - \|\mathbf{P}\mathbf{V}\|_F^2)^2 \\ &= \|(\mathbf{I}_d - \mathbf{P})\mathbf{V}\|_F^4.\end{aligned} \tag{7.61}$$

Here we used the fact that $K = \|\mathbf{V}\|_F^2 = \|\mathbf{PV}\|_F^2 + \|(\mathbf{I}_d - \mathbf{P})\mathbf{V}\|_F^2$. □

Lemma 8. *Suppose* $\mathbf{P}$ *and* $\widehat{\mathbf{P}}$ *are* $d \times d$ *projection matrices with rank* $K \leq d$. *Then* $\|\mathbf{P}\widehat{\mathbf{P}}\mathbf{P} - \mathbf{P}\|_F^2 = \sum_{j=1}^K [1 - \sigma_j(\mathbf{P}\widehat{\mathbf{P}}\mathbf{P})]^2$.

Proof. There exists $\mathbf{Q} \in \mathcal{O}_{d\times d}$ such that $\mathbf{QPQ}^T = \mathbf{D} = \begin{pmatrix} \mathbf{I}_K & \mathbf{0}_{K\times(d-K)} \\ \mathbf{0}_{(d-K)\times K} & \mathbf{0}_{(d-K)\times(d-K)} \end{pmatrix}$.

$$\|\mathbf{P}\widehat{\mathbf{P}}\mathbf{P} - \mathbf{P}\|_F^2 = \|\mathbf{Q}(\mathbf{P}\widehat{\mathbf{P}}\mathbf{P} - \mathbf{P})\mathbf{Q}^T\|_F^2 = \|\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T - \mathbf{D}\|_F^2. \tag{7.62}$$

We claim that $\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T = \begin{pmatrix} (\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T)_{[K][K]} & \mathbf{0}_{K\times(d-K)} \\ \mathbf{0}_{(d-K)\times K} & \mathbf{0}_{(d-K)\times(d-K)} \end{pmatrix}$, where $(\cdot)_{[K][K]}$ denotes the $K \times K$ submatrix in the upper left corner. This is because

$$\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T = (\mathbf{QPQ}^T)(\mathbf{Q}\widehat{\mathbf{P}}\mathbf{Q}^T)(\mathbf{QPQ}^T) = \mathbf{D}(\mathbf{Q}\widehat{\mathbf{P}}\mathbf{Q}^T)\mathbf{D}. \tag{7.63}$$

Lemma 9 yields

$$\begin{aligned} \|\mathbf{P}\widehat{\mathbf{P}}\mathbf{P} - \mathbf{P}\|_F^2 &= \|(\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T)_{[K][K]} - \mathbf{I}_K\|_F^2 = \sum_{j=1}^K \{1 - \sigma_j[(\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T)_{[K][K]}]\}^2 \\ &= \sum_{j=1}^K [1 - \sigma_j(\mathbf{QP}\widehat{\mathbf{P}}\mathbf{PQ}^T)]^2 = \sum_{j=1}^K [1 - \sigma_j(\mathbf{P}\widehat{\mathbf{P}}\mathbf{P})]^2. \end{aligned} \tag{7.64}$$

□

Lemma 9. *If* $\mathbf{A} \in \mathbf{I}_{r\times r}$ *is symmetric, then* $\|\mathbf{A} - \mathbf{I}_r\|_F^2 = \sum_{j=1}^r [1 - \sigma_j(\mathbf{A})]^2$.